ROCKWOOL Foundation Berlin
Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES 139/26

Delaying Fertility, Advancing Careers: The Lasting Consequences of Growing Up with a Safety Net

Matias Giaccobasso

www.rfberlin.com MAY 2026

Delaying Fertility, Advancing Careers: The Lasting Consequences of Growing Up with a Safety Net

Authors

Matias Giaccobasso

Reference

JEL Codes: J13; J16; J22; I38; H53

Keywords: Income support; fertility timing; labor markets; gender inequality.

Recommended Citation: Matias Giaccobasso (2026): Delaying Fertility, Advancing Careers: The Lasting Consequences of Growing Up with a Safety Net. RFBerlin Discussion Paper No. 139/26

Access

Disclaimer

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com

Delaying Fertility, Advancing Careers: The Lasting Consequences of Growing Up with a Safety Net*

Matias Giaccobasso†

VATT Institute for Economic Research & FIT

May 8, 2026

Abstract

How early-life income shapes women's life trajectories is central to understanding social mobility and gender inequality. This paper shows that income support during childhood delays motherhood, promotes early labor market participation, and ultimately leads to higher cumulative earnings and improved living standards. The analysis draws on eighteen years of administrative records and implements a regression discontinuity design that exploits an arbitrary eligibility rule in a large-scale, government-implemented cash transfer program in Uruguay. An additional USD 1,000 of income support during childhood, corresponding to a 13.3% increase in cumulative transfers, increases women's total months of formal employment by 2.5 (6.8%) and cumulative earnings by USD 1,740 (7.0%), with outcomes observed on average at age 28. These gains operate through changes in transitions to adulthood: an additional USD 1,000 of benefits increases the probability of a career-oriented transition by 3.4 percentage points (14.2%), explained by a reduction in the probability of a teen birth of 2.5 p.p. (11.6%), an increase in the probability of being employed by age 19 of 2.3 p.p. (13.4%), and an increase in tertiary education enrollment of 0.9 p.p. (8.1%). These changes place women on later fertility paths with smaller and less persistent child penalties. A Marginal Value of Public Funds analysis shows that, through increased future payroll tax revenues, the program fully pays for itself by age 37 and generates approximately USD 2.4 in government revenue per dollar transferred over the life cycle. Overall, these results show that income support during childhood can reshape women's fertility timing and early career trajectories, promoting social mobility and reducing gender inequality.

†VATT Institute for Economic Research and Finnish Centre of Excellence in Tax Systems Research (FIT). Email: matias.giaccobasso@vatt.fi.

*I am very grateful to Ricardo Perez-Truglia, Paola Giuliano, Manisha Shah, and Marcelo Bergolo for their continuous support throughout this project. I thank Natalie Bau, Sebastian Edwards, Clémence Tricaud, Nico Voigtländer, Romain Wacziarg, and Melanie Wasserman for their thoughtful comments and guidance. I also thank Sebastian Calonico, Matias Cattaneo, Raj Chetty, Javier Feinmann, Miko Hallikainen, Jarkko Harju, Hilary Hoynes, Henrik Kleven, Erzo Luttmer, Juliana Londoño-Vélez, Tuomas Matikka, Maria Sauval, Darío Tortarolo, Andrea Vigorito, and participants at the 2024 NBER Trans-Atlantic Public Economics Seminar and the 2022 NBER Public Economics Program Meeting, as well as conference participants at the 2023 PACDEV, 2023 MWIEDC, 2024 ZEW Public Finance Conference, 2024 EAYE, and 2025 ESPE. This project was supported by the Center for Global Management at the UCLA Anderson School of Management. I am also grateful for funding from the Research Council of Finland, grant no. 346253. Romina Quagliotti provided superb research assistance.

Introduction

Fertility timing plays a central role in shaping women's life trajectories. Delaying fertility, especially from early ages, has substantial labor market returns for women (Goldin and Katz, 2002; Bailey, 2006; Gallen et al., 2024), with important implications for career trajectories, social mobility, and gender inequality. As a result, governments have sought to empower women and delay marriage and childbearing through policies that explicitly target educational opportunities or reproductive choices, such as education subsidies, compulsory schooling, or access to contraceptives. At the same time, a growing literature shows that income support for poor families can also have lasting effects on women's life trajectories and labor market outcomes (Araujo and Macours, 2021; Bailey et al., 2024). Such policies have the potential to affect later-life outcomes through multiple mechanisms, including changes in financial constraints, expectations, parental investments, and early-life education and fertility decisions. However, in practice, how, when, and why money matters remain unclear (Page, 2024). This paper links these two strands of literature by showing that income support during childhood delays motherhood, shifting women toward more career-oriented transitions to adulthood, ultimately resulting in higher cumulative labor market earnings and improved living standards.

The context is a permanent, large-scale and government-implemented program: the Uruguayan *PANES/AFAM-PE*. It consists of a cash transfer that, originally, represented close to 70% of self-reported pre-program household income. It was implemented in 2005 and remains in place today as the most generous anti-poverty program in Uruguay's history, accounting for 0.4% of the Uruguayan GDP and reaching more than 10% of households. This income-support program provides immediate financial assistance to socio-economically disadvantaged households and aims to encourage medium- and long-term human capital accumulation, particularly among children. Eligibility is determined by a poverty score. Accepted households are, on paper, subject to some conditionalities such as school enrollment, attendance, and health check ups, but these were not regularly enforced until 2013.

The empirical analysis uses an exhaustive individual-level longitudinal dataset covering the universe of program applicants, built from several administrative records. This allows me to track, at the individual level, program applications and participation, as well as education, fertility, and labor market outcomes for about eighteen years. A key advantage of this longitudinal structure is that it allows me to study how income support affects outcomes dynamically rather than focusing on outcomes measured at a single age. Education, fertility, and labor market decisions are strongly correlated and interact differently over time. As emphasized by Aizer et al. (2022) and evidence from in-utero shocks and early childhood interventions (e.g., Chetty et al. 2011; Almond et al. 2018), policy effects may evolve non-linearly over the life cycle, so analyses based on snapshots of outcomes may miss important dynamic responses, such as shifts in labor market entry or delays in fertility, that ultimately shape longer-run outcomes. Following individuals across multiple life stages is therefore critical to characterize the mechanisms through which childhood income support affects life trajectories.

In this study, I focus on individuals who were younger than 18 years old when their parents first applied to the program, and at least 19 years old in December 2021. To identify the causal effects of additional income support during childhood, I use a Regression Discontinuity Design

(RDD) that exploits the sharp change in the total amount collected from *PANES/AFAM-PE* at the eligibility threshold. Because the program has been in place for more than 15 years, many households have filed multiple applications. Moreover, as the eligibility threshold became more lenient over time, many households that were initially ineligible eventually entered the program. Hence, following the approach in Jepsen et al. (2016), I use eligibility based on the score from the first application as an instrument for the total amount of childhood income support. On average, households that were ineligible at their first application collected about USD 7,600 (in 2008 PPP dollars), while eligible households collected around USD 10,900. This implies an average discontinuity of USD 3,300 (43.4%). For reference, this amount is roughly 3.3 additional years of *household formal labor* income during childhood, or about 1.4 years of *self-reported household total* income (including informal income and other transfers).

I begin by showing that, among women who were children when their parents first applied to *PANES/AFAM-PE*, additional childhood income support improves later-life formal labor market outcomes. The effects are positive but moderate for outcomes measured in the last observed year. When these women are last observed, at an average age of 28, each additional USD 1,000 of income support, equivalent to an increase of roughly 13% of cumulative transfers at the cutoff, increases last-year months worked by 3.0% and last-year earnings by 5.2%. The effects are larger and more precisely estimated when outcomes are measured cumulatively over early adulthood, providing stronger evidence of persistent gains: cumulative months employed increase by 6.8% and cumulative labor earnings by 7.0% per additional USD 1,000 received during childhood. Scaling these estimates by the eligibility-induced increase in transfers of USD 3,300 implies cumulative reduced-form effects of 22.3% for months employed and 23.2% for earnings, corresponding to roughly 8 additional months of employment and USD 5,700 in additional earnings. For men, estimated effects are smaller, and not robust across specifications.

Second, I show that women's long-term labor market improvements are driven by changes in transitions to adulthood, and that the moderate last-year effects partly reflect changes in fertility dynamics rather than a reversal of earlier gains. To summarize changes in transitions to adulthood, I use a measure of a *career-oriented* transition index, defined as not having a birth before age 19 and either holding a job or being enrolled in tertiary education at that age. Each additional 1,000 USD of income support increases the probability of a career-oriented transition by 3.4 p.p. (14.2%). Consistent effects are estimated along all dimensions: a 1,000 USD increase reduces the probability of a teen birth by 2.5 p.p. (11.6%), increases the probability of having a job before age 19 by 2.3 p.p. (13.4%), and increases the likelihood of tertiary education enrollment by 0.9 p.p. (8.1%). A dynamic analysis shows that employment and fertility effects move in opposite directions over the life cycle, with employment effects rising at the ages when fertility effects decline. By age 30, the negative effects on the probability of having given birth disappear, suggesting that women who delay motherhood catch up at later ages. This fertility catch-up may temporarily reduce labor supply, helping explain why last-year labor market effects are more moderate than cumulative effects. The dynamic analysis also shows positive effects on women's early career mobility, followed by greater stability later on, consistent with early career job exploration followed by more stable matches, potentially facilitated by delayed fertility.

Third, I explore how more career-oriented transitions translate into improved labor market outcomes, and whether this happens primarily through postponed motherhood. I begin by showing, descriptively, that additional income support has stronger effects among non-mothers than among mothers. I then focus on earnings dynamics around childbirth among women who become mothers. This analysis shows that in the Uruguayan context the relative child penalty declines with age at first birth. Furthermore, combining a stacked DiD approach with the RDD, I show that the effects of *PANES/AFAM-PE* on earnings increase over time among not-yet-mothers, while child penalties conditional on age of first birth are similar for eligible and ineligible mothers. Altogether, these results highlight two ways in which fertility postponement raises cumulative earnings: by shifting women into age-at-first-birth trajectories with lower child penalties, and by extending the period over which earnings gains accumulate prior to childbirth. Conditional on age at first birth, however, differences in the earnings response to childbirth across eligible and ineligible mothers are limited. This conditional comparison should be interpreted cautiously, since eligibility itself affects the timing of first birth and therefore the composition of women observed in each age-at-first-birth trajectory.

Beyond its main estimates, the paper provides additional evidence on why women may have changed their transitions to adulthood, focusing on secondary education outcomes and heterogeneity by age at exposure. While I document improvements in secondary education progression that point to a role for human capital accumulation, heterogeneity by age at exposure analysis also shows that effects on age at first birth and labor market entry are larger among older cohorts, who had limited scope to adjust educational attainment. This suggests that delayed motherhood cannot be explained by education decisions alone, and that other channels, such as relaxed financial constraints or foot-in-the-door mechanisms, may also play a role (Bailey et al., 2017). A notable departure from the main mechanism arises for those aged 2-5 when the program was rolled out, for whom the program primarily increases university enrollment at the expense of early labor market participation, with little evidence of changes in fertility timing. I discuss how this pattern may reflect both the contemporaneous expansion of access to contraception and abortion and the presence of dynamic complementarities, whereby early-life investments increase the returns to later career investments, highlighting how mechanisms depend on both the timing of exposure and the institutional context. Finally, I show that long-term improvements extend beyond formal labor income, as women exposed to additional childhood income support experience sizable reductions in poverty scores when interacting with the welfare system as adults.

To interpret these findings through a unified theoretical framework, I extend the career–fertility timing model in Doepke et al. (2023) introducing a childhood period in which households make schooling decisions and women face teenage fertility risk, shaping the conditions for transitioning to adulthood. In this framework, higher childhood income increases schooling investment and reduces exposure to teenage fertility risk, shifting fertility to later ages and allowing for greater early-career investment, with lasting consequences for labor market outcomes. This framework helps understanding how early-life income may reshape incentives and constraints for fertility and career investment decisions over the life cycle.

The effects documented throughout the paper have direct and important policy implications. Using the Marginal Value of Public Funds framework proposed by Hendren and Sprung-Keyser (2020), I estimate that, through increased payroll tax revenues, the *PANES/AFAM-PE* fully pays for itself already by age 37, implying an infinite MVPF. Projecting earnings gains up to retirement age, I further estimate that the program generates government revenues of approximately USD 2.4 per dollar transferred over the working life. These findings are consistent with Hendren and Sprung-Keyser (2020), who show that policies investing directly in low-income families tend to have the highest MVPFs and often pay for themselves. Beyond its potential self-financing nature, income support programs may offer additional advantages relative to policies that target specific margins more directly, such as fertility (e.g., contraceptive access) or education (e.g., schooling subsidies). By relaxing household budget constraints, childhood income support may affect a broader set of outcomes, including children's health, nutrition, and education (Fiszbein et al., 2009; Bastagli et al., 2016; Almond et al., 2018; Molina Millán et al., 2019), parental investments (Gennetian et al., 2024; Krause et al., 2025), and potentially generate positive spillovers (Egger et al., 2022).

While the results discussed so far correspond to a specific institutional context, both the empirical and theoretical results suggest that the underlying mechanisms rely on economic incentives and constraints that are fairly general and present in many contexts, including meaningful variation in fertility timing, so that there is margin for women to adjust the timing of first birth, child-related earnings penalties, and the relevance of early-career investment for subsequent labor market outcomes. At the same time, the magnitude of the effects depends on key mediating factors, such as the size and age profile of child penalties, gender norms, and labor market structure. In settings where child-related earnings penalties are small or do not vary with age, the returns to delaying fertility are limited, weakening the link between fertility timing and earnings. Similarly, in environments with more rigid gender norms or weaker labor market opportunities for women, fertility decisions may be less responsive to economic incentives or career investment may be less rewarded, reducing the scope for this mechanism to generate sizable effects.

This paper contributes to two broad literatures that are central to understanding gender inequality and social mobility. First, a large body of work documents that gender gaps in earnings arise primarily around childbirth, as women experience persistent labor market penalties following motherhood (Angelov et al., 2016; Kleven et al., 2019, 2025), despite similar earnings trajectories prior to birth (e.g., Bertrand et al. 2010; Goldin 2014; Blau and Kahn 2017). Consistent with this idea, recent work shows that delaying fertility, especially from early ages, has substantial labor market returns for women (Goldin and Katz, 2002; Bailey, 2006; Miller, 2011; Gallen et al., 2024).[1] Second, a growing literature shows that income shocks during childhood can have sizable and lasting effects on later-life outcomes (e.g., Akee et al. 2010; Almond et al. 2018; Bulman et al. 2021), including predictable, policy-driven income support and social safety net programs (e.g., Aizer et al. 2016; Manoli and Turner 2018; Bastian and Michelmore 2018; Barr et al. 2022; Bailey et al. 2024), with evidence of stronger long-run responses among women (e.g., Hoynes et al. 2016; Araujo and Macours 2021; Bastian et al. 2022; Parker and Vogl 2023; Bitler and Figinski 2024), across different contexts. However, how, when, and why money matters remains unclear (Page, 2024). This paper connects these literatures and provides an answer to these questions by showing that childhood income support shifts women toward more

[1] Similar evidence comes from settings in which women's ability to delay childbearing is restricted rather than expanded. In particular, recent work shows that abortion denial have negative lasting consequences for their educational and labor market outcomes (Miller et al., 2023; Londoño-Vélez and Saravia, 2025).

career-oriented, rather than family-oriented, transitions to adulthood, by delaying fertility, with lasting consequences for women's human capital accumulation and labor market trajectories. By highlighting this mechanism, the paper clarifies a channel through which social safety net policies shape career trajectories, gender inequality and social mobility over the long run.

Beyond its main contribution, the paper also speaks directly to the literature on the causes and consequences of teenage fertility (Hotz et al., 2005; Berthelon and Kruger, 2011; Kearney and Levine, 2012, 2014, 2015), and in particular to the role of parental income during childhood (Chetty et al., 2014). In this regard, Kearney and Levine (2012, 2014, 2015) emphasize the role of economic marginalization and limited future opportunities in shaping early childbearing decisions. From this perspective, teenage fertility reflects, at least in part, poor educational and labor market prospects rather than fertility preferences alone. Consistent with this view, the results in this paper provide complementary evidence that additional childhood income support can delay first births by changing schooling decisions and early labor market attachment. Moreover, the observed changes in secondary education progression are consistent with human capital accumulation and expected wage mechanisms (Black et al., 2008; DeCicca and Krashinsky, 2020), although strong responses by older cohorts are also consistent with income effects or foot-in-the-door explanations (Bailey et al., 2017). By documenting how fertility timing responds to changes in childhood socio-economic conditions, the paper contributes to a better understanding of the causes and consequences of teenage childbearing.

1 Institutional Background

Uruguay is a middle-high-income country with about 3.5 million inhabitants. In 2005, it had the third-highest GDP per capita in South America and ranked 53th in the Human Development Index, within the high-HDI group. Educational attainment and demographic indicators, however, had room for improvement. Lower secondary completion was 59.3%, similar to Argentina but lower than Chile (76.0%), and far from richer countries such as the United States, Sweden, or Italy. Adolescent fertility was also relatively high (60.3 per 1,000), comparable to Argentina and Costa Rica but well above Chile (50.8) and most high-income countries (e.g. Sweden: 5.7; Spain: 11.5; United States: 39.5; OECD average: 32.7).[2] Its social protection system is one of the oldest and most developed in Latin America. Family allowances were introduced in 1943 but remained restricted to formal (registered) workers. *PANES/AFAM-PE* was introduced in 2005 as a response to the early-2000s economic crisis and expanded family income support to informal workers. The program was expanded in 2007 and remains in place today.

***PANES/AFAM-PE* design.** *PANES/AFAM-PE* is a cash transfer program targeted at the poorest households in the country.[3] In the short run, the program aims to provide income support to vulnerable households and reduce poverty. In the longer run, it aims to promote human capital

[2]See Appendix A.1 for further details.

[3]*PANES* also included additional components targeted to households in the most critical conditions, such as temporary public employment programs, education, and job training. These reached a small share of beneficiaries (around 15%) and by design allocated to households very far away from the eligibility threshold. Data on participation in these components is, unfortuntately, not available. Given their limited reach and targeting, for the purpose of this study *PANES* can be treated as a cash transfer program. A more detailed description is provided in Appendix A.

accumulation and social mobility, especially among children. Its cost is about 0.4% of GDP, similar to major programs like *PROGRESA-Oportunidades* (Mexico) and *Bolsa Familia* (Brazil), and close to what the United States spends on family benefits (0.67% of GDP).[4] It consists of two phases, *PANES* and *AFAM-PE*, which I treat as a broader income support policy: *PANES/AFAM-PE*. The first phase was in place from 2005 to 2007, and the second phase started immediately afterward and remains in place today. The goals and structure of the program stayed largely the same but some changes were introduced aimed at increasing its coverage. In practice, most *PANES* participants were automatically enrolled in *AFAM-PE*, along with some households that were initially rejected as eligibility rules became more lenient. Broadly speaking, *AFAM-PE* can be thought of as an expansion of *PANES* with minor changes in the payment structure.

During PANES, the benefit consisted of a base payment of USD 133 and a complement between USD 29 and USD 78 for households with underage children (about 70% of participating households).[5] The base amount alone represents roughly 63.9% of self-reported household income at the time of application, and 78.1%. when the supplementary payments are included. These values, however, should be interpreted with caution, as households may have faced incentives to under-report income.[6] Under *AFAM-PE*, the transfer amount was determined as a base payment of USD 57 per underage child plus a USD 24 supplementary payment per child aged 12–18 enrolled in secondary school, both subject to a 0.6 equivalence scale.

Between 2005 and 2007, more than 180,000 different households (17.6% of all households in the country) applied to *PANES/AFAM-PE*. Eligibility was determined based on a proxy-means test. Households were visited by program officials who surveyed their socio-economic and material conditions. This information was used to compute a poverty score (z), which represents the predicted probability of having a household income in the bottom 20% of the distribution.[7] Households with a poverty score z above a certain region-specific threshold are eligible to participate, while households with a score below the threshold are deemed ineligible.[8] After being accepted, participant households were supposed to satisfy school attendance, regular health check-ups, and monthly per-capita income requirements. However, the program did not rigorously enforce the education and health conditions until June, 2013.

Existing Studies. There are several studies that document how *PANES/AFAM-PE* affected parental outcomes. For instance, when looking at adults' labor market decisions, Bergolo and Cruces (2021) finds that the program led to a 13% reduction in formal employment, explained both by changes in labor supply and formality decisions. There is also a series of studies that

[4]https://data.oecd.org/socialexp/family-benefits-public-spending.htm

[5]Values are expressed in January 2008 PPP dollars using CPI and PPP conversion factors. In local currency, the USD 133 corresponded to UYU 1,360.

[6]As an alternative reference, in 2005, the monthly transfer for a household with four or more children was equivalent to one minimum wage.

[7]The variables used to calculate the score included the overall quality of the building, the number of people living in the household, the number of rooms, the presence of underage children, average years of education, and type of employment, among others. More details about how the poverty score was computed can be found in Appendix A and in Manacorda et al. (2011) or in Amarante et al. (2016).

[8]Formally, households could be rejected for other reasons than the score. For instance, the program also required households to pass an income test at the moment of application. Only households with per capita income below approximately US$131 a month were eligible and were subsequently visited by personnel from the Ministry for Social Development. However, the income condition was not particularly binding as it only disqualified around 10 percent of the initial applicants (Amarante et al., 2013).

focus on adults' fertility, birth outcomes, and household structures. These studies show that the program did not affect older women's fertility (Parada, 2024), but improved birth outcomes conditional on giving birth (Amarante et al., 2016). Regarding household structure and intra-household decisions, the program led to more stability in pre-program marital status (Parada, 2023), and did not change actual power structures within the household, although it improved women's self-perceived agency in decisions about household expenditure (Bergolo and Galván, 2018).[9] Focusing on the children, Bloomfield and Cabrera (2026) show that the program led to overall moderate-to-weak effects on educational attainment, but somewhat stronger effects for some children exposed to the program when they were 0-5. Similarly, Amarante et al. (2013) find null effects of *PANES/AFAM-PE* on child labor and school attendance, although these results relied on very short-term survey data collected during the first two years of the program.

2 Data, Sample, and Outcomes of Interest

2.1 Data Sources, Treatment, and Outcome Variables

***PANES/AFAM-PE* records: Application and participation variables.** These records are used to construct all application- and participation-related variables. They were provided by the Ministry of Social Development, and correspond to all *PANES/AFAM-PE* applications and payments between April 2005 and December 2023. Application data was collected during the survey aimed at assessing eligibility conditions and includes information on city, application date, poverty score, resolution, housing conditions (e.g., materials, access to sanitation, appliances, etc.), education, employment status, and other personal characteristics for each household member.[10] Participation records include information on household monthly payments, which are used to compute the total amount of cash transferred by *PANES/AFAM-PE*.

Employer-employee matched records: Labor market outcomes. These were provided by the Ministry of Labor and Social Security and cover the universe of formal employment in Uruguay from 2005 to 2023, including all registered private- and public-sector workers as well as the self-employed. The records contain monthly job-level information such as start and end dates, earnings, and firm characteristics. The main outcomes built from these data are: last-year and cumulative months worked and earnings. For the dynamic analysis, I also construct age-by-age employment indicators as well as the age of the first first formal job, defined as the first job spell of at least four consecutive months, to exclude temporary jobs.[11] All outcome variables are constructed using only events after the first application.

One caveat of using administrative records is that they allow precise measurement of *formal* employment and earnings but do not capture *informal* work. For this reason, the analysis focuses on formal labor market outcomes, which are themselves meaningful indicators of job quality,

[9]Other studies have focused on preferences and attitudes. Manacorda et al. (2011) finds that the program increased support for the government and trust in the President and institutions implementing the program, while Nicolau (2023) find that the program might have slightly increased self-reported feelings associated with stigma.

[10]Appendix B contains a more detailed description of the application and participation data, as well as descriptive statistics for different samples of interest.

[11]Appendix C provides summary statistics for each of the outcome variables used in the analysis.

including stability, access to social insurance, and attachment to the contributory labor market.[12]

Birth Records: Fertility outcomes. Birth records were provided by the Ministy of Public Health and are used to measure fertility outcomes. They contain information on the universe of births in Uruguay between 2003 and 2021, including birth date, birth weight, gestation weeks, and mother and father identification information. Information on fathers should be taken with a grain of salt since this information is only collected starting in 2010, when electronic records started and, even in the best years, about 50% of fathers' identifiers are missing. The two main fertility outcomes are age at first birth and age-specific indicators for being a parent by a given age. As secondary outcomes, I also use the total number of children ever born and an indicator for being a parent at the last observed date. In this case, the post-treatment period is defined as starting seven months after the application date.[13]

Secondary and tertiary education records: Education outcomes. Secondary schooling records were provided by the National Administration of Public Education and cover lower and upper secondary education (grades 7 to 12) from 2004 to 2024, including information on enrollment and progression/achievement. Tertiary education records come from the *Universidad de la República*, the country's largest university with roughly 80 percent of national enrollment, and include the universe of students ever enrolled in any major between 2005 and 2020. Coverage is universal for public schools, but current data does not capture enrollment in vocational or private schools, so responses along these margins are not observed.[14] Importantly, the (free) public education system is likely the relevant choice set for the population of interest, as private institutions typically offer few grants and charge relatively high tuition.[15] For secondary education, the main outcomes are ever enrolling, years enrolled, the maximum grade attained, and an indicator for ever enrolling in 12th grade.[16] For tertiary education, the main outcome is an indicator for ever enrolling in the Universidad de la República after the first application.[17]

2.2 Sample of Interest and Descriptive Statistics

The *full sample* is comprised by 420,383 individuals who were younger than 18 years old when their parents first applied to *PANES/AFAM-PE*, and at least 19 years old by December 2021, the last year of birth records. In addition, I further restrict the sample to simplify the interpretation of results and exclude households with outlier applications. First, to ensure that comparisons

[12] According to official statistics, informal employment was 21.3% of total employment in Uruguay in 2023.

[13] Because all outcomes are constructed after the first application, fertility outcomes refer to births observed after the first application. I refer to these as *first births* for simplicity, since pre-application births are very rare in the analysis sample (0.78%)

[14] Access to data on vocational schools has been authorized but, as of this writing, the data have not yet been provided. Data for private institutions are only available starting in 2023.

[15] For instance, Ramírez Leira (2021) shows that more than 95% of individuals in the first income quintile were enrolled in public institutions in 2017.

[16] Unfortunately, Uruguay does not have centralized records on secondary school completion, which would be the natural outcome to measure secondary education achievement.

[17] One caveat is that tertiary education records are observed only through 2020. As a result, although, as described below, the final sample includes individuals aged 2–17 when *PANES/AFAM-PE* was rolled out, estimates of the effects on university enrollment by age at exposure for the youngest group mainly capture responses among those who were 3–5 years old in March 2005.

are made around a single, region-specific threshold, the analysis is restricted to individuals from households that first applied to *PANES/AFAM-PE* between 2005 and 2007, when the eligibility threshold remained constant. This accounts for 66.5% of individuals in the *full sample*. Second, I drop 24,551 individuals in the bottom and top 5% of the poverty score distribution, as observations in the extreme tails of the running variable tend to add noise when computing optimal bandwidths or including polynomial functions of the score in the RD specifications. Finally, I exclude 6,182 individuals with outlier application behavior.[18] After applying these filters, the final sample consists of 248,693 individuals, which I refer to as the *main sample*.

Individuals in the *main sample* are on average 9.8 years old at the time of their household's first application and about 28 years old in December 2023. On average, they appear in 3.2 application forms over the period, including separate counts for the *PANES* and *AFAM-PE* phases, as well as, in some cases, later applications as household heads or partners in newly formed households. The sample is roughly balanced by gender (48.5% are women), and households have on average 4.9 members, with an average member age of 21.0 years and 45.3% of them being single-parent households. Household heads report employment rates of 64.6%, average 6.3 years of education, which is slightly above completed primary school, and labor income of about USD 121.3, which is well below the poverty line (approximately USD 350). When restricting attention to individuals close to the eligibility threshold, most characteristics remain extremely similar except, by construction, for variables related to the poverty score.[19] For instance, individuals who are close to the cutoff are concentrated around zero and have lower first-application acceptance rate (57.4% vs. 76.0%), as the overall distribution of the poverty score is heavily skewed toward the eligible region. Other noticeable differences are the lower share of applications submitted in the capital city (15.5% versus 27.1%), and a slightly larger share of single-parent households (52.9% versus 45.3% for the full support sample).[20]

3 Empirical Strategy

3.1 Research Design

The discontinuity in eligibility status at the threshold provides quasi-random variation to identify the causal effects of *PANES/AFAM-PE* using a Regression Discontinuity Design (RDD). Intuitively, under perfect compliance and a continuity assumption, (local) average treatment effects of the program can be obtained by comparing the regression functions of the outcome of interest at both sides of the threshold (Imbens and Lemieux, 2008; Hahn et al., 2001).[21]

From the perspective of the program administrators, i.e., treating each application form

[18]This includes 2,328 individuals in more than eight application forms, 2,112 from households with more than seven applications, and 1,742 listed as household heads or partners at the time of the first application.

[19]For simplicity, the bandwidth used for this comparison is 0.02, which corresponds to the maximum optimal bandwidth reported in Table 2, the main set of labor market results.

[20]Further details and a full set of descriptive statistics are provided in Appendix C, including a description of the universe of applicants, a comparison between the *main* and *full sample*, and additional statistics on the main outcomes of interest. Overall, the two samples look very similar overall; the main differences relate to application behavior and are explained by the restrictions defining the main sample.

[21]Formally, let Y be any of the outcomes of interest. Under perfect compliance, the key identification assumption in RDD is that the potential outcome regression functions, $\mathbb{E}[Y(1) \mid Z = z]$ and $\mathbb{E}[Y(0) \mid Z = z]$, are continuous at $z = 0$. Under this assumption, $\mathbb{E}[Y(1) - Y(0) \mid Z = 0] = \lim_{z \downarrow 0} \mathbb{E}[Y \mid Z = z] - \lim_{z \uparrow 0} \mathbb{E}[Y \mid Z = z]$.

as the unit of observation, the eligibility rule was followed very closely: roughly 80% of score-ineligible applications were rejected, and about 90% of score-eligible applications were accepted. However, for a small number of ineligible forms extremely close to the threshold (within 0.0015, or 0.15 p.p.), acceptance rates are abnormally high. This is most likely due to precision issues in the raw data, which in some cases were recorded with only four decimal places. Throughout the analysis I will then exclude application forms within 0.15 p.p. of the threshold.[22]

From a household perspective, exposure to the program does not depend on a single application form, but on the full history of applications and participation. Since *PANES/AFAM-PE* has been in place since 2005, families could have re-applied multiple times, and it is therefore necessary to track household application and participation histories rather than individual forms. This creates two main challenges for the RDD: (i) deciding which application score to use as the running variable, and (ii) avoiding potential endogenous sorting around the eligibility threshold induced by re-application decisions. To address these challenges, I follow the approach in Jepsen et al. (2016), who propose using the score from the first application as the running variable in RDD settings with multiple applications. Intuitively, this is the less manipulable score and avoids endogeneity associated with selection into re-applications. Hence, the baseline RDD strategy can be characterized as follows:

Exogenous variable: eligibility based on the score of the first application form. I define the instrument D^{1st} as a binary variable that takes the value of 1 if the score obtained in the first application is above the region-specific threshold and 0 otherwise. I define the *first* application for an individual as the earliest application among all households in which the individual ever appears as a member. For example, if a household applies to *PANES* with two parents and an older child, and later appears in *AFAM-PE* with the same members plus a younger child born after the first application, the relevant first application for the younger child is the earlier *PANES* application of the household.[23]

Endogenous variable: amount collected from *PANES/AFAM-PE* during childhood. I define the baseline treatment variable (T) as the total amount of cash transferred by *PANES/AFAM-PE* to the individual's household before age 18, expressed in thousands of 2008 PPP USD, winsorized at the 95th percentile.[24] For outcomes measured before age 18, I construct analogous treatment variables for each corresponding age. In addition, I also report estimates based on two alternative endogenous variables: (i) months enrolled in the program, capped at 192 and (ii) ever participating.

I estimate τ_{FRD}, i.e., the effect of an additional USD 1,000 of childhood income support on outcome Y, using the following specification:

$$Y_i = \mu + \tau_{FRD} T_i + \beta_1 Z_i^{1st} + \beta_2 Z_i^{1st} T_i + \mathbf{X_i} + \mathbf{\Lambda} + u_i \tag{1}$$

[22] Appendix D provides additional details on how the eligibility rule was applied. For the main results, I also present robustness checks that vary the donut radius. Estimates are not sensitive to this choice.

[23] Appendix B describes how I create household links across administrative records.

[24] For children who lived in more than one beneficiary household during childhood, I sum income support across households. All results remain unchanged when using only the amount transferred to the first household. For conciseness, I do not report estimates based on this alternative endogenous variable in the paper.

where Y_i is the outcome of interest for individual i, (Z_i^{1st}) is the centered value of the poverty score obtained in the first application, T_i is the endogenous treatment variable, $\mathbf{X_i}$ is a series of individual-level baseline characteristics (i.e., those reported in Table 1, except for the predicted eligibility) and $\mathbf{\Lambda}$ represents year of birth, year of application, and region FE. Treatment variation in T_i is instrumented by eligibility D_i^{1st}, and the corresponding first-stage relationship is given by:

$$T_i = \alpha + \delta D_i^{1st} + \gamma_1 Z_i^{1st} + \gamma_2 Z_i^{1st} D_i^{1st} + \mathbf{X_i} + \mathbf{\Lambda} + \epsilon_i \tag{2}$$

Following Imbens and Lemieux (2008), estimates are based local linear regressions fitted separately to each side of the threshold, and p-values are based on robust bias-corrected inference, following Calonico et al. (2014). The RDD bandwidth is defined optimally following the data-driven approach by Calonico et al. (2019) with selection of bandwidth by optimization of Mean Squared Error (MSERD) and a triangular kernel function. In all cases, standard errors are clustered at the household level.

In addition, I also report figures and estimates corresponding to reduced-form effects, i.e., τ_{SRD}, for the main otucomes of interest based on the following sharp RDD specification:

$$Y_i = \mu + \tau_{SRD} D_i^{1st} + \beta_1 Z_i^{1st} + \beta_2 Z_i^{1st} D_i^{1st} + \mathbf{X_i} + \mathbf{\Lambda} + u_i \tag{3}$$

Compared to sharp RDDs, fuzzy RDDs require the additional identifying assumption of monotonicity (Imbens and Lemieux, 2008; Cattaneo et al., 2019). In this setting, monotonicity requires that crossing the eligibility threshold in the first application weakly increases cumulative transfers received before age 18 for all individuals in the sample (i.e., the no-defiers assumption). This assumption is plausible given that initially ineligible households can only enter the program at later ages, which mechanically limits their potential exposure before age 18 and makes it unlikely that early eligibility would reduce cumulative transfers. Additional evidence consistent with this assumption is presented in the next section.

Validity of the RDD Design and Interpretation of Estimates

Figure 1 summarizes a series of empirical tests used to validate the RDD design. Panel (a) reports the average amount of income support collected from *PANES/AFAM-PE* during childhood as a function of the score obtained in the first application (Z_i^{1st}). The figure is restricted to the optimal MSE bandwidth (1.6 p.p.) based on Calonico et al. (2019) and uses 15 quantile-spaced bins on each side of the threshold. The figure shows a sharp jump in benefits at the eligibility cutoff: individuals from households whose first-application score was just below the eligibility threshold (henceforth, *ineligible* individuals) received on average about 7,600 USD in transfers, whereas those just above (henceforth, *eligible*) collected about 10,900 USD (in 2008 values). This yields a discontinuity of roughly 3,271 USD (p-value < 0.001), or close to a 43.4% increase in childhood income support. For reference, this difference is approximately equivalent to 3.3 additional years of household formal labor income during childhood, or about 1.4 years of self-reported household total income (including informal income and other transfers such as pensions, retirement benefits, etc.). Appendix D shows that this discontinuity also represents about 20.3 months of exposure (32.7%) and a 18.4 p.p. (24.4%) increase in the probability of ever participating in *PANES/AFAM-PE*.[25]

In addition, following the approach in Rose and Shem-Tov (2021), Panel (b) shows that the first stage illustrated in Panel (a) average differences in treatment intensity coming from all across the distribution. In this case, each bin represents the local change in the probability of collecting $T \geq t$ from *PANES/AFAM-PE* during the whole period.[26] Differences are statistically significant for all values reported between USD 5,000 and USD 25,000, suggesting that the proposed instrument comprises substantial variation in intensity of exposure. Importantly, as discussed in Section 3.1, the additional monotonicity assumption required for the fuzzy RDD approach implies that the instrument, in this case crossing the eligibility threshold, should not reduce treatment for any individual. As noted in Rose and Shem-Tov (2021) and Angrist and Imbens (1995), estimates such as those reported in Panel (b) can be thought of an indirect validation of the monotonicity assumption, as eligibility increases exposure to treatment all across the distribution, showing no visible violations to this assumption.

Panel (c) in Figure 1 shows that eligibility for *PANES/AFAM-PE* substantially increased the financial resources available to children during childhood. The figure reports year-by-year sharp-RDD estimates for three income measures around the time of the first application: (i) total family income (formal earnings of all household members plus transfers), (ii) total adult income (formal earnings of adult members plus transfers), and (iii) household *PANES/AFAM-PE* income.[27] The figure shows a large jump in transfers at the time of application for eligible children, which fades out after 2–3 years, consistent with the 20-month discontinuity in program exposure discussed above.[28] The patterns for total adult and total family income are also informative. The smaller increase in total adult income relative to transfers likely reflects reductions in adults' formal labor earnings as a response to additional income support, consistent with Bergolo and Cruces (2021) who also find evidence of such reductions and further show that these are explained by real labor supply responses and formal/informal earnings shifting by equal proportions. Hence, the true change in total adult income, i.e., also including informal earnings, lies between the estimated change in total adult income (based on formal earnings and transfers) and the change in transfer income. At the same time, effects on total family income are larger than those on adult income, with a gap that increases over time, consistent with positive effects on earnings of (baseline) non-adult household members as they enter the labor market, anticipating some of the results discussed later in Section 4. Overall, these results point to a strong net increase in total household resources, indicating that eligible children grew up in households with substantially higher income due to the policy.

[25] Appendix D also provides additional descriptives on the distribution of months exposed, amounts collected, the full distribution of z, and the first-stage relationship using the entire support of z. It additionally reports a series of placebo tests showing that the first stage appears only at the true cutoff and that estimates are not sensitive to bandwidth choice, donut radius, or other technical decisions in the RDD estimation, such as the kernel function, polynomial order, or the bandwidth selection algorithm.

[26] For instance, the bin at $t = 10$ shows how much more likely it is for eligible individuals to receive USD 10,000 or more in *PANES/AFAM-PE* benefits compared to ineligible.

[27] Ideally, I would measure income available to all household members at each point in time, but household composition is only observed at the time of application. Income measures are therefore constructed using a fixed baseline household roster.

[28] Additional evidence in Appendix D shows that this fade-out is driven by ineligible households entering the program as eligibility rules became more lenient. The small increase in transfers before application reflects retroactive payments accrued from April 2005.

Finally, panel (d) in Figure 1 and Table 1 provide indirect tests of the RDD identification assumptions. First, Panel (d) reports a formal test of continuity of the running variable following McCrary (2008) and Cattaneo et al. (2018) and shows no evidence of discontinuity in the running variable at the cutoff (p-value = 0.384). Second, Table 1 shows that there are no systematic discontinuities in baseline characteristics at the threshold, which is reassuring about the causal interpretation of the RDD estimates. The table reports continuity tests for a series of baseline characteristics measured in the application form, all of which are included as control variables in the main specification for increased statistical power.[29] The first row provides a useful summary test using predicted first-time eligibility based on observable characteristics as the outcome variable. The absence of a jump in predicted eligibility ($\tau_{SRD} = -0.002$, p-value = 0.740) suggests that baseline characteristics behave smoothly around the cutoff.[30]

Interpretation of RDD Estimates. The RDD approach used next identifies the causal effect of changes in childhood income support using variation in treatment intensity induced by crossing the *PANES/AFAM-PE* eligibility threshold at first application. As panel (a) in Figure 1 and additional results in Appendix D show, differences in cumulative transfers around the cutoff may arise from several margins, including the timing of entry into the program, whether individuals ever receive transfers, and the duration of exposure conditional on entry. Panel (b) in Figure 1 further shows that these differences reflect shifts across the entire distribution of transfers. For convenience, τ_{FRD} summarizes the effect of this change in exposure in a single point estimate expressed in dollar-equivalent terms. However, when interpreting this estimate, one should keep in mind that it reflects effects induced by this bundled variation rather than a specific marginal change in benefits associated with a given part of the transfers distribution while holding the timing of transfers fixed. Similarly, while the reduced-form effect τ_{SRD} corresponds to the effect of crossing the eligibility threshold directly, without relying on any scaling assumption, it still reflects variation arising from multiple margins.

4 Results

4.1 Effects of PANES/AFAM-PE on Labor Market Outcomes

In this section, I present the main estimates on how additional childhood income support from *PANES/AFAM-PE* affected individuals' labor market outcomes as adults. Figures 2 and 4 present the reduced-form estimates for women and men separately. Panels (a) and (b) report last-year outcomes, while panels (c) and (d) show cumulative outcomes over the full follow-up period. Table 2 complements these figures by reporting the corresponding 2SLS estimates from the fuzzy RDD specification in Equation (1), together with the reduced-form and first-stage coefficients from Equations (3) and (2), respectively.

[29] The intercepts reported in Table 1 are obtained from the baseline RDD specification in Equation (3), replacing the outcome Y_i with each baseline characteristic and including the same fixed effects $\mathbf{\Lambda}$ used in the main analysis, but excluding the additional covariates $\mathbf{X}_i$ because these are the outcomes of interest in the continuity tests. Therefore, they need not coincide with the descriptive statistics reported in Appendix C.

[30] The reduced form figure for the discontinuity in predicted eligibility is reported in Appendix D.

Effects on Women's Labor Market Outcomes. Focusing on women, panels (a) and (b) of Figure 2 show positive, although moderate, discontinuities around the threshold for last-year months worked and last-year earnings. For months worked, the (reduced-form) estimated discontinuity is 0.5 months and statistically significant at the 10% level (p-value = 0.074). Relative to a baseline of 5.0 months in the control group, this corresponds to an increase of about 9.8%. The results for last-year earnings are similar. The discontinuity at the threshold is USD 654 and statistically significant at the 5% level (p-value = 0.025), which represents a 17.0% increase relative to a baseline of USD 3,850 in the control group. When looking at the 2SLS estimates, columns (1) and (2) of Table 2 show that an additional USD 1,000 of childhood income support increases women's last-year months employed by 0.15 (3.0%) and last-year earnings by USD 199 (5.2%). Both effects are statistically significant at the usual levels (p-values = 0.079 and 0.029).

The discontinuities are much sharper when considering the full early-adulthood period. Panels (c) and (d) of Figure 2 present reduced-form evidence both for cumulative total months employed and total earnings. For total months worked, the estimated discontinuity is 8.3 months, representing an increase of 22.3% relative to the ineligible group. For cumulative earnings, the discontinuity is USD 5,738, corresponding to a 23.2% increase. In both cases, estimates are strongly statistically significant at the 1% level, with p-values < 0.001 and $= 0.002$, respectively. The 2SLS estimates, reported in columns (3) and (4) of Table 2, show that an additional USD 1,000 of income support increases cumulative months worked by 2.5 (6.8%) and cumulative earnings by USD 1,740 (7.0%), with both effects being strongly statistically significant (p-values < 0.001 and $= 0.002$, respectively).

The magnitude of these effects is large. The reduced-form increase of 8.3 months corresponds to about 22.3% of the baseline mean and roughly the 30th percentile of the outcome distribution among ineligible women near the cutoff. Under a stronger linearity assumption, extrapolating the 2.5-month 2SLS estimate to a move from zero to median exposure (USD 7,600) implies an increase of about 19 months, close to the median of the distribution. Cumulative earnings gains are also sizable: the reduced-form effect alone accounts for roughly half of the observed gender gap in cumulative earnings (about USD 10,600), while extrapolation to median exposure would eliminate the gap entirely. Although these extrapolations rely on strong assumptions and should be interpreted cautiously, they reinforce the idea that the effects of additional childhood income support are economically meaningful.

Additional labor market outcomes are reported in Appendix E. First, consistent with the evidence for last-year months worked and last-year earnings, the appendix shows positive but imprecisely estimated effects on last-year employment and wages. In particular, the 2SLS effect on the probability of being employed in July 2023 is 1.4 p.p. (3.4%, p-value = 0.057), while estimates for log wages in July 2023 and log cumulative earnings over the last year (conditional on positive earnings) are also positive (1% and 2.8%) but not statistically significant (p-values = 0.182 and 0.181, respectively). Importantly, because log-outcomes are conditional on employment, they may also reflect changes in the composition of the employed sample, making their interpretation more nuanced. Second, effects are stronger when focusing on more stable measures of employment attachment. For example, the 2SLS effect on the probability of being employed for *all* 12 months of the last year is 2.1 p.p. (8.1%, p-value = 0.007). I explore these patterns further in Section 4.2, where I analyze the effects of additional childhood income support on job mobility and later-life labor market stability. Third, I examine the characteristics of firms in which women are employed in July 2023, including firm size, median wage, total wage bill, and firm age, and find no statistically significant effects. However, these estimates may also reflect changes in employee composition, as the intervention may bring marginal workers into employment who sort into different types of firms. To address this, I also report estimates of unconditional employment by sector and show that employment gains are driven primarily by retail trade, and to a lesser extent by administrative services and manufacturing.[31]

Taken together, the results discussed so far suggest positive effects on labor market outcomes, with larger and more precisely estimated effects when outcomes are measured cumulatively over early adulthood. However, when interpreting last-year outcomes, two additional considerations are relevant. First, last-year outcomes are measured using a single year of data and are therefore more exposed to transitory fluctuations in employment and earnings, which may reduce statistical precision relative to cumulative outcomes. The second consideration is more conceptual and relates to the dynamics of treatment effects over the life cycle. Last-year outcomes capture labor market attachment at the latest observed age. Because treatment effects may change over time, last-year labor market effects reflect the realization of those dynamics at older ages, rather than the full trajectory of labor market responses. I return to this in Section 4.2, where I examine fertility and labor market effects age by age and show that the smaller last-year effects are mostly driven by the specific age dynamics of the effects.

Robustness and Sensitivity. A natural concern in RDD settings is whether these cumulative effects reflect a genuine discontinuity at the eligibility threshold. To address this concern, I conduct an extensive set of robustness and sensitivity exercises. Figure 3 reports these tests for women's cumulative earnings as a leading example. First, panel (a) shows the results of a placebo cutoff exercise. In particular, re-estimating the same reduced-form specification at a dense grid of placebo thresholds yields a distribution of estimates centered around zero, with the estimate at the true cutoff standing out as uniquely large in absolute value. Panels (b) through (d) show that the estimated discontinuity is highly robust to a wide range of modeling choices. In panel (b), the 2SLS effect remains positive and statistically significant at the 10% level for all manually selected bandwidths between 0.01 and 0.1. Panel (c) shows that the estimated effect is consistently positive and confidence intervals rarely fall below 0 across more than 60 alternative RDD specifications that vary kernel functions, polynomial degrees, and bandwidth selectors. Finally, panel (d) shows that varying the donut radius from 0 to 0.003 does not affect the results. Appendix E reports analogous tests for all other outcomes in Figure 2. For cumulative months worked, the results are identical to those for cumulative earnings. For last-year outcomes, the direction of the effects is stable across specifications, but statistical significance declines with wider bandwidths, explained mostly by the smaller point estimates. The appendix also shows that results are robust to excluding additional controls ($\mathbf{X}$) and fixed effects ($\mathbf{\Lambda}$), to using alternative definitions of the endogenous variable, including months of participation and an indicator for ever receiving *PANES/AFAM-PE*, and to remove winsorizing from both the dependent variable and the outcome variables..

[31] More precisely, I define an indicator variable that takes the value of 1 if someone is employed in a given sector, and 0 otherwise, including both employment in a different sector and unemployment.

Effects on Men's Labor Market Outcomes. When looking at men, the evidence is considerably weaker overall. Panels (a) and (b) of Figure 4 show positive discontinuities for last-year outcomes, similar in sign to those for women. However, while the magnitude for last-year months worked is similar to that observed for women, the effect on last-year earnings is substantially smaller for men (9.9% vs. 17.0%). Furthermore, Appendix E shows that these effects attenuate quickly with bandwidth choice and are more sensitive to RDD specifications than the corresponding effects for women. For instance, estimates for last-year months worked and earnings become statistically insignificant around a bandwidth of 0.04 and converge to zero for wider bandwidths. More importantly, cumulative outcomes in Panels (c) and (d) do not show any evidence of discontinuities. Estimated effects on cumulative months worked (2.0, or 4.1%) and cumulative earnings (USD 817, or 2.3%) are much smaller than for women (22.3% and 23.2%) and are not statistically significant (p-values = 0.228 and 0.555). Robustness checks further show that these estimates can turn negative under alternative specifications, and even become statistically significant for wider bandwidths. Appendix E shows that this pattern of null effects extends to a broader set of outcomes.

Overall, the results discussed in this section show that additional childhood income support leads to large and robust cumulative gains for women, with last-year effects being somewhat more modest and imprecise, and little evidence of any sizable effects for men. These patterns raise two questions. First, why did the program generate substantial improvements for women? Second, why are these gains not as strongly reflected in outcomes measured in the last observed year? The next section explores these questions by examining how individuals exposed to additional income support transitioned into adulthood, the age-dynamics of the effects, and additional job mobility patterns.

4.2 Transitions to Adulthood and Labor Market Dynamics

Career-Oriented Index. Next, I explore how children who benefited from *PANES/AFAM-PE* changed their transitions to adulthood, summarizing these transitions into a single index that captures *career-oriented* rather than *parenthood-oriented* pathways.[32] This index captures behavioral responses along fertility, education, and labor market margins. I define someone as taking a career-oriented path if two conditions hold: (i) they did not have a child before age 19, and (ii) they either had their first relatively stable job (defined as an employment spell of at least four consecutive months) before age 19 or are enrolled in university.[33]

The use of the term career-oriented" is inspired by the distinction between a "job" and a "career," where careers involve stronger and more continuous labor force attachment, generating greater incentives to invest in human capital, both through formal schooling and on-the-job training (Goldin, 2004, 2006, 2021). In this paper, the concept is used more broadly to contrast transitions oriented toward additional education or labor market attachment with transitions

[32] The concept of *transition to adulthood* describes a period in which individuals make a set of closely related education, labor market, and fertility decisions that characterize the move from childhood to adult roles. In the sociology and psychology literature, this period typically spans the late teenage years through the twenties and is increasingly viewed as distinct from both adolescence and full adulthood (Arnett, 2000; Settersten Jr et al., 2008).

[33] I use age 19 because it roughly marks the end of compulsory schooling, is a common cutoff in the teenage pregnancy literature, and is observed for all individuals in the sample.

centered around early family formation. This definition is not meant to be exhaustive, but it provides a transparent way of capturing whether individuals enter adulthood primarily investing in early-career human capital (i.e., on-the-job experience or formal education) or through early family formation.

Overall, additional childhood income support increases the likelihood of career-oriented transitions among women. Panel (a) of Figure 5 shows a statistically significant increase of 10.9 p.p. at the eligibility threshold (p-value < 0.001), relative to a baseline of 23.7 p.p., while Table 3 shows that each additional USD 1,000 increases this probability by 3.4 p.p. (14.2%). Importantly, these effects are not driven by a single margin, but are present across all dimensions. Reduced-form estimates show that teen births decline by 8.2 p.p. (p-value < 0.001), relative to 21.7%, while formal employment before age 19 increases by 7.4 p.p. (p-value < 0.001), relative to 17.1%, and university enrollment rises by 2.9 p.p. (p-value = 0.048). The corresponding 2SLS estimates show that each additional USD 1,000 reduces teen births by 2.5 p.p. (11.6%), increases employment before age 19 by 2.3 p.p. (13.4%), and raises university enrollment by 0.9 p.p. (8.1%, p-value = 0.053). As a complementary approach, Table 4 examines the timing of key transitions using age-of-event outcomes, together with the corresponding extensive-margin measures. The reduced-form results show that additional childhood income support delays age at first birth by 0.5 years (from a baseline of 20.0, p-value = 0.001) and brings forward the age at first stable job by 0.8 years (from 20.7, p-value < 0.001), with corresponding 2SLS estimates of 0.16 and -0.26. At the same time, the table documents no significant effects on the probability of becoming a mother or ever being employed, alleviating concerns about selection into these outcomes. In sum, these results show that women who grew up with additional childhood income support are more likely to choose career-oriented transitions rather than paths associated with early family formation and limited labor market participation.

As for labor market outcomes, the magnitude of these effects is also large. For instance, the reduced-form effect on the teenage birth rate is around 40%, roughly two thirds of the total decline in teenage birth rates observed since 2005 (from 60 to 23 births per 1,000 women aged 15–19). Income support also appears relatively more effective in reducing teenage fertility than other contemporaneous policy changes, such as abortion legalization (Cabella and Velázquez, 2022) or expanded access to LARCs (Ceni et al., 2021), though this comparison should be cautious as *PANES/AFAM-PE* targeted poorer women, who face higher baseline fertility risk.

Regarding men, Tables 3 and 4 show that, in general, the program does not seem to have significantly affected their transition to adulthood. However, it is important to note that column (6) in Table 4 shows a -0.8-year reduced-form effect on the age at which men become fathers. While this estimate is statistically significant (p-value = 0.006), it should be taken with with a grain of salt due to important limitations in the availability of fathers' identifying information, as discussed in Section 2. For conciseness, full estimates are reported in Appendix F.

Dynamic Effects. To explore more directly how the timing of these key events changed, I exploit the longitudinal nature of the data and estimate age-by-age effects. Figure 6 summarize these dynamics for two outcomes of interest.[34] For women's fertility-related behavior, the

[34] Appendix F includes the corresponding full tables for all estimates shown in these figures both for the 2SLS and reduced-form estimates.

outcome is defined as having given birth by a given age γ, which allows me to track cumulative differences in the timing of motherhood across ages. For labor-market outcomes, the outcome variable is being employed at age γ, capturing how labor market participation effects evolve over time. In each figure, the y-axis reports the 2SLS estimate and the x-axis the age at which the outcome is measured.[35] Because the *main sample* includes individuals observed at different ages, I use age-specific samples to exploit as much information as possible. This means that effects at age γ are estimated using women aged γ or older. While this could raise concerns that the documented patterns reflect changes in sample composition rather than true dynamics, Appendix F shows that the results are robust to using a balanced sample.

Panel (a) of Figure 6 shows a clear age pattern of women's fertility responses. Effects are negative and concentrated in the late teens and early twenties, with larger magnitudes observed between ages 15 and 19. The effects remain large and statistically significant until around age 25, and then fade out, becoming close to zero and statistically insignificant by age 30. Two conclusions follow from this pattern. First, the postponement documented in Tables 3 and 4 is mainly driven by a reduction in births that would have occurred during the teenage years and early twenties. Second, the program shifts the timing of births rather than overall fertility, at least at the extensive margin. Women who delay fertility begin to catch up around age 25, and by age 30 there are no significant effects on the probability of having children. Naturally, postponing the first birth may eventually lead to lower completed fertility simply because women have less time to reach their desired number of children. Appendix F shows a reduced-form decline of 0.1 in total births (-11.4%, p-value = 0.025), although this should be interpreted as suggestive since women in the sample have not yet completed their reproductive cycles.

A similar but oppositely signed pattern is observed when looking at women's dynamic effects on employment. The overall shape of these dynamics almost mirrors the image of the fertility effects: the period of strongest fertility reductions corresponds to the period of strongest employment gains, and both sets of effects start to fade out at around the same age. These patterns are consistent with childbearing and labor-market participation being competing activities for women. In fact, these results help interpreting the relatively moderate effects on last-year earnings discussed in Section 4.1: by the end of the observation period, many women who had delayed fertility earlier begin to catch up in childbearing, reducing contemporaneous labor supply, at least temporarily.

For men, and consistent with the reduced-form and 2SLS results discussed in Section 4.1, the dynamic estimates reported in Panel (b) of Figure 6 do not show statistically or economically significant effects in either fertility or employment outcomes. This reinforces the conclusion that the program's effects on transitions to adulthood are driven almost entirely by women.

Labor Market Mobility. Finally, I look at job mobility dynamics more directly. Table 5 reports estimates for: (i) the number of firms in which a woman held her main position, (ii) the probability of entering the labor market and remaining in the same firm until last observed, (iii) indicators of low (1–2 employers throughout the period), moderate (3–4 employers), and high

[35] For instance, the coefficient shown for $\gamma = 25$ corresponds to the effect of an additional USD 1,000 in childhood income support on the probability that an individual has had their first child at or before age 25, or in the probability of being employed at age 25.

mobility (5+), and (iv) tenure-related outcomes among women employed in July 2023. Overall, estimates show that women receiving additional income support are somewhat more mobile. For instance, the reduced-form effect of *PANES/AFAM-PE* eligibility on the probability of staying in the same firm they started is -4.1 p.p. (20.8%, p-value = 0.027), driven mainly by low-mobility women shifting into moderate-mobility paths (three to four employers). Importantly, the program does not increase the probability of very high mobility (five or more employers). At the same time, Columns (6)–(8) show that women who received additional support are more likely to hold more stable jobs when last observed, conditional on being employed. There is a positive reduced-form effect of 0.11 log points (p-value = 0.021) on log tenure in the current job and an increase of 7.9 p.p. (18.0%, p-value = 0.006) in the probability of having at least two years of tenure. Taken together, these results suggest that treated women may have engaged in more job mobility early on and eventually settle into more stable positions.

Overall, the evidence discussed in this section suggests that the long-term improvements in labor market outcomes arise because women shift toward more career-oriented transitions to adulthood, mainly by avoiding early births and entering the labor market earlier. These changes also manifest in greater (moderate) job mobility early on and more stable jobs when last observed, likely facilitated by the additional time and flexibility for early-career job search, which in turn leads to better and more stable longer-term matches. More broadly, these dynamics reflect a reorganization of women's transitions to adulthood, without meaningful changes in underlying preferences for motherhood or formal labor market participation. While these patterns do not by themselves establish a causal link, they point to a strong connection between delayed fertility and improved early-career labor market outcomes. Next, I investigate this link more directly.

4.3 Labor Market Dynamics Around Childbirth

A first, although somewhat naive, way to explore the link between delayed fertility and improved early-career labor market outcomes is to compare effects between mothers and non-mothers. Intuitively, if women's early-career investment and labor market gains are driven by delayed motherhood, one would expect stronger improvements for non-mothers. Table 6 shows that not having children is positively correlated with stronger labor market effects. Across all four outcomes, estimated effects for non-mothers are larger than for mothers. For instance, reduced-form estimates for non-mothers show increases of about 8.9 cumulative months employed and USD 5,931 in cumulative earnings (25.7% and 23.8%), while for mothers the corresponding effects are 5.9 months and USD 3,702 (14.9% and 14.8%).

While these findings are consistent with labor market improvements being driven, at least in part, by delayed or foregone motherhood, women who remain childless may differ systematically from those who become mothers, and may therefore not be the most appropriate comparison group. I therefore focus on women who have children and examine how changes in fertility timing change labor market dynamics around childbirth. Conditional on becoming a mother, changes in the timing of first birth can affect labor market dynamics in two ways: by shifting women across age-at-first-birth trajectories with different child penalties, and by affecting earnings dynamics around childbirth conditional on age at first birth. I next provide exploratory evidence on both channels to assess whether the observed labor market gains are primarily driven by changes in fertility timing or by differences in earnings dynamics conditional on age at first birth.

I begin by studying differences across age-at-first-birth trajectories using a stacked DiD approach (Wing et al., 2024). Using a stacked DiD rather than the most typical event-study approach is justified by the fact that mothers who give birth at different ages follow different labor-market trajectories, and the penalty from childbirth varies substantially by age (e.g., Adda et al. 2017; Gallen et al. 2024). Pooling them into a single event-study implicitly assumes that mothers who give birth at different ages are comparable and that childbirth has homogeneous effects across ages, which is unlikely and can bias standard TWFE estimates (Roth et al., 2023; De Chaisemartin and d'Haultfoeuille, 2023).

Following Melentyeva and Riedel (2023), I construct a stacked dataset of separate sub-events for each age at first birth s and compare women who give birth at age s to not-yet-mothers who will give birth within the next three years. This comparison aims to ensure that mothers and not-yet-mothers are on similar labor market paths but do not differ critically in the timing of motherhood. A key requirement of the stacked DiD design is that each age-at-first-birth group be observed over the full pre- and post-birth window. This requirement determines the age restrictions used in the analysis, which is therefore limited to women who gave birth between ages 21 and 27. The resulting stacked dataset removes variation in treatment timing within each sub-event by construction, allowing me to estimate how child penalties vary with the age at first birth using the following specification:

$$Y_{isa} = \sum_{\substack{l=-5 \\ l\neq -2}}^{3} \beta_{sl} \left(B_{is} \times \mathbb{1}[a - s = l]\right) \; + \; \gamma_{is} \; + \; \lambda_{sa} \; + \; \kappa_t \; + \; \varepsilon_{isa}, \tag{4}$$

where Y_{isa} denotes the outcome of interest for woman i in stack s observed at age a, and B_{is} is an indicator equal to one for women who give birth at age s and zero for not-yet-mothers who will give birth at ages $s+1$ through $s+3$. Control units are included in the stack only while they remain not-yet-mothers, so they do not contribute post-birth outcomes and the number of controls declines at longer event times. The term $\mathbb{1}[a - s = l]$ is an event-time indicator equal to one when an observation is l years away from the focal age s. Thus, $l = 0$ corresponds to age s, $l = -1$ to one year before that age, and so on. For treated women, this coincides with years relative to childbirth, while for control women it corresponds to years relative to the comparison age s. The coefficients β_{sl} are the objects of interest and trace the evolution of outcomes around childbirth separately for each age at first birth s, relative to the omitted event time $l = -2$. The fixed effects γ_{is} and λ_{sa} are individual-by-stack and age-by-stack fixed effects, respectively, which flexibly absorb differences in earnings levels and life-cycle profiles within each stack, and κ_t denotes calendar-year fixed effects. Standard errors are clustered at the individual level.

Figure 7 summarizes the main results from the stacked DiD approach and provides evidence on heterogeneous earnings trajectories around childbirth by age-at-first-birth. In both panels, the outcome of interest is annual earnings expressed as a percent of the within-stack average earnings two years before birth. Panel (a) reports separate point estimates by age at first birth, while panel (b) aggregates ages into two groups (21–23 and 24–27) and reports 95%

confidence intervals.[36] The results show that relative earnings losses are substantially larger for women who give birth at younger ages, and decline monotonically with maternal age. Recovery dynamics also differ sharply across trajectories. For women who become mothers earlier, the earnings penalty continues to increase several years after birth, while for women who give birth later the evolution is, at worst, flat.[37] Overall, these results show that delaying motherhood shifts women into age-at-first-birth trajectories characterized by substantially lower relative child penalties, highlighting fertility timing as a key channel through which additional childhood income support can translate into improved labor-market outcomes.[38]

To examine the second channel, i.e., whether additional income support affects earnings dynamics around childbirth conditional on age at first birth, I combine the stacked DiD design with the quasi-random variation generated by the *PANES/AFAM-PE* eligibility threshold. This specification holds fertility timing fixed and tests whether, among women who give birth at the same age, exposure to additional income support changes earnings responses to childbirth. This can be interpreted as a triple-difference-type design that compares earnings dynamics around childbirth for eligible and ineligible mothers, while using not-yet-mothers as a counterfactual to net out the effects of additional childhood income support on earnings trajectories prior to childbirth. Note again that, in this stacked DiD setup, event time is defined relative to age s, so that periods with $l \geq 0$ correspond to post-birth outcomes for treated women but still reflect pre-birth outcomes for not-yet-mothers. This approach therefore aims to separate the effects of *PANES/AFAM-PE* on general earnings trajectories from its effects on earnings responses to childbirth itself. However, one important caveat is that, because age at first birth is itself affected by the program, this conditional comparison can reflect compositional differences within age-at-first-birth groups. In particular, eligible women giving birth at a given age may include women who would have given birth earlier in the absence of the program, and therefore differ from ineligible women giving birth at that same age.

The following specification combines Equations (2) and (4) by fully interacting *PANES/AFAM-PE* eligibility status and score with motherhood status and event time.

$$\begin{aligned} Y_{isa} = & \sum_{\substack{l=-5 \\ l\neq -2}}^{3} \alpha_l \, D_i^{1st} \, \mathbb{1}[a-s=l] + \sum_{\substack{l=-5 \\ l\neq -2}}^{3} \gamma_l \, B_{is} \, \mathbb{1}[a-s=l] \\ & + \sum_{\substack{l=-5 \\ l\neq -2}}^{3} \beta_l \, B_{is} D_i^{1st} \, \mathbb{1}[a-s=l] + \sum_{\substack{l=-5 \\ l\neq -2}}^{3} f_l(Z_i^{1st}, D_i^{1st}, B_{is}) \, \mathbb{1}[a-s=l] \\ & + \gamma_{is} + \lambda_{sa} + \kappa_t + \varepsilon_{isa}. \end{aligned} \tag{5}$$

where α_l captures how eligibility shifts earnings for not-yet-mothers relative to ineligible not-yet-mothers at the same age and within the same stack, γ_l captures the earnings response to childbirth for ineligible mothers relative to ineligible not-yet-mothers, and β_l measures

[36]To aggregate effects across age-at-birth groups, I follow Melentyeva and Riedel (2023) and weight the group-specific estimates by their sample shares (Sun and Abraham, 2021; Wing et al., 2024).

[37]These patterns are consistent with previous findings in different contexts, such as Sweden, Germany, and the United States, which also document a negative relationship between age at first birth and the career or earnings costs of children (Miller, 2011; Adda et al., 2017; Melentyeva and Riedel, 2023; Gallen et al., 2024).

[38]Appendix F reports additional results using absolute earnings and extensive-margin outcomes.

differences in the earnings response to childbirth between eligible and ineligible mothers, conditional on age at first birth. The term $f_l(Z_i^{1st}, D_i^{1st}, B_{is})$ is a flexible function of the running variable that varies by event time l and interacts Z_i^{1st} with eligibility and motherhood status.[39] For simplicity, the sample is restricted to individuals within 0.05 of the eligibility cutoff and the regression is estimated on the pooled stacked sample and weighted so that each stack contributes in proportion to its treated share. Hence, the coefficients α_l, γ_l, and β_l represent average effects across the stacked age-at-first-birth groups. Standard errors are clustered at the individual level.

Estimates from this model are depicted in Figure 8. First, the pre-birth estimates support the identification assumption as barely eligible and ineligible women show similar earnings dynamics prior to childbirth within age-at-first-birth groups (i.e., parallel pre-trends). Second, the figure show that not-yet-mothers who received additional income support during childhood experience positive and increasing earnings gains over event time. These effects are captured by the α_l coefficients and reflect the impact of *PANES/AFAM-PE* on earnings trajectories that would continue to accumulate in the absence of childbirth. Third, once these general earnings effects are netted out, there is little evidence of substantial differences in the earnings response to childbirth experienced by eligible and ineligible mothers at a given age at first birth.[40]

Overall, the results reported in this section show that by shifting births to later ages, additional income support moves women into age-at-first-birth trajectories characterized by substantially lower and less persistent child penalties. At the same time, delaying motherhood extends the period over which eligible women experience positive earnings gains prior to childbirth, allowing these gains to accumulate over a longer horizon. Together, these two forces generate large cumulative earnings advantages relative to women who did not receive additional income support. Importantly, however, once fertility timing is held fixed (i.e., conditional on age at first birth), there is little evidence that the program drastically changes the earnings response to childbirth itself, although this comparison may partly reflect compositional differences within age-at-first-birth groups.

4.4 Additional Evidence on Mechanisms and Heterogeneity

Effects on Secondary Education Outcomes A natural follow-up question is which factors could explain shifts in women's fertility choices. One hypothesis is that changes in educational choices may affect fertility decisions directly. In particular, the literature has emphasized both incapacitation and human capital accumulation mechanisms (Black et al., 2008). Incapacitation operates through increased time spent in controlled environments, where teenagers are less likely to engage in risky behavior. Human capital accumulation, in turn, raises expected returns in the formal labor market, increasing the opportunity cost of early motherhood

Using data on secondary education enrollment, I examine a series of outcomes capturing participation and progression in middle and high school. Overall, Table 7 shows that additional income support during childhood is associated with progression to higher grades, but not with changes in overall enrollment. In particular, eligibility increases the probability of reaching 12th grade and the maximum grade attained, while effects on ever being enrolled and on years of

[39]In practice, this includes local-linear terms in Z_i^{1st}, as well as their interactions with D_i^{1st} and B_{is}, all interacted with the event-time indicators.

[40]Appendix F shows similar findings using a TWFE combined with an RDD approach in the unstacked data.

enrollment are close to zero and statistically insignificant. While these results are exploratory, they suggest that additional income support may have shifted educational trajectories toward higher attainment rather than increasing enrollment per se. This pattern is more consistent with a human capital channel, with higher expected returns to schooling explaining delayed fertility, than with a pure incapacitation mechanism operating through increased time spent in school.

Heterogeneity by Age at Exposure A growing literature shows that earlier childhood exposure to additional resources can have larger long-run effects than exposure later in childhood. (e.g., Bailey et al. 2024). At the same time, recent work has shown that interventions delivered at critical moments, such as adolescence, can also generate sizable effects (Page, 2024). Estimates reported in Figure 9 explore whether the effects documented in this paper vary by age at *PANES/AFAM-PE* rollout, which I use as a proxy for the timing of exposure as the timing of application itself may be endogenous. Importantly, all estimates are obtained from the same 2SLS specification and are scaled per USD 1,000 of additional income support. Thus, differences across cohorts do not reflect that some cohorts received mechanically larger transfers than others; rather, they indicate that the estimated effect of a given amount of additional income support differs by age at exposure.

Two main conclusions can be drawn from this analysis. First, women who were exposed to *PANES/AFAM-PE* after age 5, drive the main results discussed so far. These cohorts show positive effects on cumulative earnings and months worked, fertility delays, and earlier labor market participation. In terms of the age gradient, the effects on cumulative labor-market outcomes appear roughly similar across cohorts. In contrast, effects on age at first birth and first job appear larger for cohorts exposed at older ages. Since these cohorts had limited scope to adjust educational attainment, this pattern suggests that changes in secondary schooling alone cannot explain the observed fertility responses. One alternative explanation is that cash transfers are more effective at relaxing financial constraints when delivered at high-risk moments, in line with cash-on-hand or foot-in-the-door-type mechanisms (Bailey et al., 2017).

The youngest cohort, i.e., ages 2–5 when *PANES/AFAM-PE* was rolled out, displays a different response: effects on cumulative labor-market outcomes are mostly negative, effects on age at first birth and age at first job are close to zero, and effects on university enrollment are large and positive. As will be discussed in Section 5, this pattern is consistent with a dynamic complementarities interpretation where earlier exposure to additional resources leads to greater human capital accumulation prior to the transition to adulthood, which in turn raises the returns to career-oriented investments and makes them disproportionately more attractive (Heckman, 2006; Cunha and Heckman, 2007; Cunha et al., 2010). The negative labor-market effects for this cohort align with this mechanism, as increased university enrollment, a time-intensive form of career investment, likely comes at the expense of labor-market participation at younger ages.

The absence of fertility responses for this cohort is less intuitive, since tertiary education is often viewed as a substitute for early fertility. A possible explanation is that, by the time this cohort reached adolescence, Uruguay had expanded access to fertility control through long-acting reversible contraceptives and legalized abortion, both of which reduced teenage fertility (Ceni et al., 2021; Cabella and Velázquez, 2022). These reforms may have allowed untreated women to delay fertility as well, narrowing fertility-timing differences by childhood income.[41]

Exploring Fertility and Education as Mediating Factors Next, I provide some descriptive and exploratory evidence on to what extent labor market effects can be accounted for by changes in fertility timing and secondary education. First, I examine heterogeneity in cumulative labor-market outcomes across three groups: women who had an early birth, women who did not have an early birth and attained low education, and women who did not have an early birth and attained high education. Figure 10 suggests that avoiding an early birth is necessary to observe positive effects on labor-market outcomes, while education mainly affects their magnitude among women without an early birth. Second, Table 8 reports a mediation-style exercise that adds controls for age at first birth and secondary education to the main specification in Equation (1). The estimates are substantially attenuated once age-at-first-birth controls are included, suggesting that fertility timing is an important mediator. Among women who gave birth, these controls reduce the estimated earnings effects by about 40 percent, while secondary education explains little additional variation. Importantly, both of these results should be interpreted as a descriptive or accounting exercise rather than as a formal causal test of mechanisms, as they rely on controlling for variables that were also affected by the program.

Downstream Effects on Adult Welfare Participation and Living Standards Finally, I examine how additional childhood income support affected living standards beyond formal earnings, once beneficiaries formed their own households. Table 9 replicates the baseline analysis for outcomes related to subsequent welfare participation, including application and acceptance status, age at first application, number of payments received, and the poverty score at the time of application. The results show at most small negative effects on participation-related outcomes, with only a marginally significant reduction of 4 months in the number of payments received (18.5%; p-value = 0.054). However, there is clear and robust evidence of improvements in poverty scores, with a reduced-form effect of -0.034 (15.0%; p-value = 0.003). Because the poverty score captures broader household conditions and is less volatile than formal earnings, this result provides a complementary measure of material well-being. These findings suggest that while childhood income support did not substantially alter welfare participation in adulthood, a finding that is likely explained by the expanded coverage of the program over time, it did lead to meaningful improvements in women's living standards beyond formal labor-market outcomes.

5 Discussion

5.1 Childhood Income, Career Investments, and Fertility Timing

This section presents a stylized framework that helps organize and interpret the results discussed so far by linking observed patterns in fertility, career investment, and earnings to some key mechanisms. The framework builds on the career-fertility timing model in Doepke et al. (2023),

[41] One caveat of this analysis is that fertility outcomes for the youngest cohort are observed only up to relatively early ages, so effects on fertility timing at older ages cannot yet be fully assessed. However, the mean age at first birth in this population is around 20 and, as shown in the main results, the strongest effects are concentrated between ages 15 and 19. Hence, this alternative explanation is unlikely to be the main reason for the null effects in this group.

which formalizes the idea that the opportunity cost of having children at young ages includes not only foregone wages, but also lower human capital accumulation through missed career opportunities, while delaying fertility is associated to higher infertility risk at later ages. The model described next preserves these core features, but adds an additional period to show how early-life income shapes the conditions under which women transition to adulthood and its implications for fertility timing, career investment, and longer run labor market outcomes. Full derivations are included in Appendix G.

Childhood Income, Schooling, and Early Fertility Risk. Time is divided into three stages: childhood and adolescence (period 0), the transition to adulthood (period 1), and adulthood (period 2). Period 0 is the main extension relative to Doepke et al. (2023) and is introduced to show how the conditions women face at the beginning of adulthood are shaped by childhood income. For simplicity, I assume that decisions in period 0 are made within the household. Thus, this period only matters as it shapes the conditions under which women make their transition-to-adulthood decisions, abstracting from a full intergenerational optimization problem.

In period 0, households allocate resources between childhood consumption c_0 and schooling investment s_0. Preferences are separable in consumption, and schooling, which is modeled as a normal good, following a large literature on education and human capital investment (e.g., Todd and Wolpin 2006, 2008; Keane and Wolpin 2010).[42] Schooling has a monetary cost p. If an early fertility shock occurs, the marginal cost of schooling increases by a factor $(1+\varphi)$, with $\varphi > 0$. This captures tighter time and resource constraints associated with early fertility, such as childcare expenditures or forgone earnings of other household members providing childcare support (Klepinger et al., 1999).[43]

Importantly, I do not model teenage/childhood fertility as a choice variable. Instead, it occurs with probability $\Pr(n_0 = 1 \mid y) = r(y)$, with $r'(y) < 0$. This reduced-form assumption summarizes a range of mechanisms documented in the literature. Higher-income households may face fewer financial barriers to contraception and reproductive health services (Bailey, 2006; Kearney and Levine, 2009; Bailey et al., 2023), provide greater parental supervision and safer environments, or raise the perceived returns to schooling and future labor market participation, thus increasing the opportunity cost of an unintended early pregnancy (Black et al., 2008). More broadly, it is consistent with evidence that economic conditions and future opportunities play a central role in shaping teenage fertility beyond preferences alone (Kearney and Levine, 2012, 2014).[44] Rather than modeling these mechanisms explicitly, the function $r(y)$ summarizes how childhood economic conditions translate into exposure to teenage childbearing.

Conditional on childhood income, optimal schooling investment is given by:

$$s_0^*(y, n_0) = \frac{\delta}{1+\delta} \cdot \frac{y}{p(1+\varphi n_0)} \tag{6}$$

[42] It is beyond the scope of this paper to discuss the non-pecuniary benefits of schooling or if it should be considered a (normal) consumption good. Oreopoulos (2011) and MacLeod and Urquiola (2019) provide in-depth reviews about the status of this discussion in the literature.

[43] As emphasized by Keane and Wolpin (2010), monetary and non-pecuniary costs are observationally equivalent in this context, so modeling early fertility through the budget constraint provides a transparent representation without loss of generality.

[44] One could obtain the same results by modeling early fertility as the result of preferences, constraints, and circumstances. The probabilistic approach adopted here provides a reduced-form representation.

where $\delta > 0$ captures the relative weight placed on schooling in household preferences. Together with the probability of the early fertility shock, this expression shows that, in this setting, childhood income affects schooling through two channels: (i) by relaxing resource constraints, and (ii) by reducing exposure to the high-cost-of-schooling state associated with early fertility.

Career Investment and Fertility Timing. Transition and adulthood decisions follow the same structure as in Doepke et al. (2023): women live for two periods and decide over fertility and career investment. During the transition period, women may invest time in their careers (e_1), which raises future wages. Fertility choices include having children in period 1, delaying childbearing to period 2, or remaining childless. Women derive utility v from having a child, which is independent of the timing of birth. Delaying fertility allows for greater career investment but carries a higher risk of infertility ($1 - \pi$). Children impose time costs (ϕ) that reduce available hours for working and career investment. Importantly, career investment should be interpreted broadly as any use of time that raises future wages but is not fully remunerated at the current stage. This includes formal training and education, but also on-the-job learning, job search for better matches, accepting lower wages in exchange for future career opportunities, or firm requirements that are incompatible with having children and are required for certain careers or to improve prospects of career progression (Goldin, 2014; Wasserman, 2023a,b).[45]

By incorporating period 0 choices and shocks, the model changes the incentives and constraints shaping women's career and fertility decisions, departing from Doepke et al. (2023) in two ways. First, childhood schooling affects labor market returns both through entry wages and through wages later in life. Entry wages satisfy $w_1 = w_1(s_0)$ with $w_1'(s_0) > 0$, while second-period wages take the form $w_2 = w_1(s_0)\kappa(s_0)e_1^{\gamma(s_0)}$. Hence, additional schooling raises the productivity scale of career investment, $\kappa(s_0)$, and also the elasticity of period-2 wages with respect to career investment, $\gamma(s_0)$. I assume $\kappa'(s_0) > 0$ and $\gamma'(s_0) > 0$, capturing dynamic complementarities between early-life human capital and later career investments (Heckman, 2006; Cunha and Heckman, 2007; Cunha et al., 2010). Intuitively, higher schooling may give access to career paths where investments have higher payoffs and steeper returns, making time-intensive career investments more attractive and increasing optimal career investment.

Second, the realization of a teenage fertility shock constrains fertility choices at the transition to adulthood. If $n_0 = 1$, the woman enters period 1 having already had a birth before adulthood, so fertility timing is no longer a choice. In this case, I set $n_1 = 1$, meaning that the period-1 time cost of childrearing applies to women who had a teenage birth before adulthood. If $n_0 = 0$, fertility timing remains endogenous. In this sense, teenage childbearing acts as a pre-transition constraint, whereas fertility in early adulthood reflects an active timing decision.

Assuming the same time-separable utility as in Doepke et al. (2023), optimal decisions are characterized by a career investment rule and a set of fertility thresholds.

$$e_1^*(s_0, n_1) = \frac{\gamma(s_0)}{1 + \gamma(s_0)}, (1 - \phi n_1), \tag{7}$$

$$v \geq \bar{v}_1 \equiv -\log(1 - \phi), \tag{8}$$

[45]These amenities include long or inflexible hours, face time, and client- or meeting-intensive schedules.

$$v \geq \bar{v}_2(s_0) \equiv \frac{(1+\gamma(s_0)) - \pi}{1-\pi}\big[-\log(1-\phi)\big], \tag{9}$$

Regarding career decisions, conditional on a fertility plan (n_1, n_2), optimal career investment is shaped by time constraints and returns: early fertility ($n_1 = 1$) reduces available time, while higher returns $\gamma(s_0)$ increases the incentive to allocate time toward career investment. Regarding fertility choices, women with $n_0 = 0$ are characterized by two thresholds in the value of children, v, where $\bar{v}_1$ separates delayed fertility from childlessness and depends only on the time cost of children, while $\bar{v}_2(s_0)$ separates early from delayed fertility and depends on returns to career investment and infertility risk. These decision rules follow the same general structure as those in Doepke et al. (2023). The key difference, however, is that early-life income affects the conditions under which they operate. By shaping schooling and exposure to teenage fertility shocks, childhood income changes entry wages, returns to career investment, and the set of fertility options at the transition to adulthood.

Comparative Statics. The comparative statics described next summarize how an increase in childhood income, such as that induced by *PANES/AFAM-PE*, affects schooling, fertility timing, career investment, and earnings through the channels described above.

At childhood, higher income increases schooling by relaxing resource constraints and reducing exposure to early fertility shocks, which raise the effective cost of schooling. It also lowers the probability of teenage childbearing, reducing the share of women who enter adulthood having already had a child. At the transition to adulthood, higher childhood income affects both fertility and career investment decisions. Fertility responses operate mainly on the timing margin: early fertility declines due to both reduced exposure to unintended teenage births and higher opportunity costs of early childbearing, shifting women from early to delayed fertility without directly affecting desired fertility. At the same time, career investment increases unambiguously, as higher schooling raises the returns to investment and delayed fertility relaxes time constraints.

Predictions for earnings are more nuanced. Hence, it is useful to distinguish between underlying earnings capacity, which corresponds to wages in the model and is driven by schooling and career investment, and realized earnings at a given age, which also depend on labor supply and fertility timing. During the transition to adulthood, higher childhood income raises wages through schooling but also strengthens incentives for career investment, tightening contemporaneous time constraints. In later periods, higher schooling and accumulated investment also raise wages, but realized earnings depend on fertility timing, as delayed childbearing may shift time costs to higher-wage periods. As a result, effects on earnings at specific ages are generally ambiguous, as wage gains may be offset by increased career investment or the time costs of children. By contrast, over the life cycle, higher schooling and career investment should generate persistent gains in wages, so cumulative earnings are expected to rise even when age-specific earnings may be temporarily attenuated.

Linking Theory to Evidence. The empirical findings of this paper are consistent with these predictions. First, estimates for secondary education and teenage fertility support the idea that childhood income shapes initial conditions through both human capital accumulation and exposure to early fertility risk. Section 4.4 shows that additional income support increases secondary education progression and maximum grade attainment, with little effect on enrollment margins, suggesting that households use additional resources to improve educational trajectories among children already in school rather than to extend time enrolled. The same section also documents fertility responses among cohorts exposed to *PANES/AFAM-PE* at later ages, when scope to adjust formal schooling is limited. This evidence is consistent with the direct income–early fertility channel emphasized in the model ($r(y)$).

The framework also helps interpret the dynamic patterns observed in fertility outcomes at later ages. Empirically, Section 4.2 shows that childhood income support delays the transition to motherhood rather than reducing the overall probability of becoming a mother. This pattern mirrors the model's emphasis on fertility timing.

The evidence is also consistent with changes in childhood income reshaping career investment decisions during the transition to adulthood. In the model, career investment captures time-intensive activities that raise future wages but are not fully remunerated contemporaneously. One way this may be reflected in the data is through job mobility and employment dynamics. Section 4.2 shows that additional childhood income is associated with higher job mobility early in life and more stable jobs at later ages, consistent with greater search and experimentation that lead to better and more persistent matches. The timing of employment and fertility responses further supports this interpretation: employment effects do not materialize immediately and peak in the mid-twenties, just before fertility effects begin to fade. This pattern suggests a reallocation of time toward career-building activities whose returns materialize gradually but attenuate as women catch up in fertility. Alternatively, one could also look at additional education beyond childhood, e.g. tertiary education. In this case, the empirical evidence supporting this mechanism comes from cohorts exposed to income support at very young ages (2–5), who exhibit substantial increases in university enrollment. This pattern is consistent with the idea that larger improvements in early-life schooling, through full childhood exposure, may lead to disproportionately higher career investment due to dynamic complementarities, including investment in very time-intensive activities. Importantly, this type of investment, such as tertiary education, is more likely to crowd out a larger share of time from paid labor market activities than job exploration or match experimentation, implying weaker or delayed earnings gains early in adulthood, despite substantial returns when accumulated over the life cycle.

Finally, the labor-market effects documented in Section 4.1 are consistent with the model's prediction that age-specific earnings effects may be more ambiguous, while cumulative earnings unambiguously increase over the life cycle. Furthermore, the differences by age of exposure illustrate how this ambiguity can arise. For cohorts exposed at later ages, career investment is more likely to operate through less time-intensive activities such as job search and experimentation, which are compatible with positive earnings and employment gains early in adulthood, with attenuation occurring later as delayed fertility is eventually realized. By contrast, for cohorts exposed at younger ages, time-intensive investments, such as college enrollment, may crowd out paid labor market activity early in adulthood, leading to slower earnings growth or even negative effects during early adulthood. From a longer-run perspective, the key empirical finding of the paper is that childhood income support leads to higher cumulative labor-market earnings, even when year-specific effects are muted or negative at certain stages of the life cycle.

5.2 Generalizability of the Results

This section discusses the extent to which the results in this paper can be generalized to other settings. A useful starting point is to clarify the conditions under which the mechanisms described above are most likely to operate. First, childhood income must affect early-life conditions such as schooling and teenage fertility risk. Second, women should have margin to adjust fertility timing, childbearing must have persistent effects on earnings, and early-career investment must be relevant for future labor market outcomes. The Uruguayan context satisfies these conditions as secondary education completion rates are relatively low and teenage fertility remains moderate to high, leaving margin for changes along both dimensions, as described in Section 1. Moreover, as documented throughout the empirical analysis, there is also meaningful variation in fertility timing, childbearing has sizable and age-dependent earnings penalties, and early adulthood is a period in which career investment matters for subsequent outcomes, as reflected by positive returns to education in the labor market.

While these economic incentives and constraints are fairly common across settings, the magnitude of the effects may depend on the scope that childhood income has in shaping women's conditions at the transition to adulthood and on how strongly these initial conditions translate into fertility timing and career investment decisions. The conceptual framework developed in the previous section helps organize these forces and highlights three key mediators: the size and shape of child penalties (e.g., ϕ), gender norms (captured by v and the implied thresholds), and labor market structure (through the specific forms of w_1, w_2, κ, and γ). Differences across countries in these dimensions should therefore be taken into account when extrapolating the results to other settings.

First, consider the role of the child penalty. The effect of postponing fertility on labor market earnings depends on the magnitude and age profile of child penalties. Uruguay's child penalty lies near the middle of the international distribution documented by Kleven et al. (2025). Hence, the estimated effects are therefore likely to be smaller in settings with limited or weakly age-dependent penalties, and larger where these are substantial and vary strongly with age at first birth. Second, consider the role of gender norms. Gender norms shape both the magnitude of the child penalty and the responsiveness of fertility timing to economic incentives. According to the Gender Social Norms Index (GSNI), Uruguay ranks 17th in the world in terms of the lowest share of individuals exhibiting at least one gender bias (UNDP, 2023).[46] In more traditional settings, where gender norms imply stronger and more persistent labor market penalties for motherhood, postponing fertility may generate larger gains. However, fertility timing may also be less responsive to economic incentives if social norms constrain family formation decisions. Conversely, in more egalitarian contexts, fertility choices may be more responsive, but smaller child penalties may limit the gains from delaying fertility. Third, consider labor market structure, particularly female labor force participation and informality. Uruguay is relatively formal and gender-integrated by regional standards (Galván et al., 2024). In settings with higher informality or weaker female labor market attachment, the returns to career investment may be lower or more uncertain, reducing the incentives to delay fertility. As a result, the link between initial conditions and life-cycle outcomes may be attenuated, particularly in terms of observed labor

[46] Field et al. (2026) show Uruguay at the bottom of the distribution in agreement with the statement "*When jobs are scarce, men should have more right to a job than women*", based on data from the World Values Survey (Wave 7).

market earnings. Moreover, because informal labor market activity is not fully observed, higher informality makes the interpretation of earnings effects more nuanced, as part of the response may happen outside the formal sector and go unobserved, while shifts into formal employment may inflate measured earnings effects.

Finally, it is also useful to clarify what population and environment the results speak to (List, 2020). The estimates are identified from households close to the eligibility threshold of a large anti-poverty program, and therefore are most informative about relatively poor households, though not those in extreme poverty. At the same time, several features of the setting strengthen the external validity of the findings. Outcomes are measured from administrative data, allowing individuals to be followed over long periods with no attrition. Moreover, the policy was implemented as a permanent, large-scale government program, affecting more than 10% of households. As a result, decisions are made in a high-stakes, real-world environment, rather than in an experimental or artificial setting. Taken together, these features suggest that the results reflect behavior under realistic conditions, while still benefiting from credible quasi-experimental variation. To further clarify how these results relate to other settings, the following section compares them to related studies, with a focus on the magnitude of the effects.

5.3 Comparison with related literature and Magnitude of the Effects

The direction of the effects documented in this paper is consistent with a growing literature showing that additional family income can generate lasting improvements in children's adult outcomes, particularly when transfers are targeted to low-income families and are regular and predictable (Page, 2024), and align with recent evidence from the United States (Aizer et al., 2016; Bastian and Michelmore, 2018; Bitler and Figinski, 2024; Bailey et al., 2024) and from Latin America (Araujo and Macours, 2021; Attanasio et al., 2021; Parker and Vogl, 2023). Importantly, the stronger responses for women are also in line with recent empirical work in the United States (Bastian et al., 2022; Hoynes et al., 2016; Bitler and Figinski, 2024) and in other contexts such as PROGRESA (Araujo and Macours, 2021; Parker and Vogl, 2023).[47]

The estimated effects are large relative to other contexts. *PANES/AFAM-PE* appears to generate stronger earnings responses than most safety-net policies studied in the United States. The 7.0% increase in cumulative earnings per additional USD 1,000 reported in Table 2 exceeds the 2.2% reduced-form effect associated with additional USD 1,000 EITC income support in Bastian and Michelmore (2018) and the 1% benchmark suggested in Page (2024). It is closer to, though still larger than, the 3–4% effects found for Food Stamps in Hoynes et al. (2016) and Bitler and Figinski (2024). These differences may reflect both program generosity, as Food Stamps and *PANES/AFAM-PE* represent large shares of household resources (Bitler and Figinski, 2024), and differences in beneficiary characteristics, as income support tends to generate stronger responses among more disadvantaged households (Page, 2024).[48] At the same time, cross-context comparisons should be interpreted cautiously, as a given dollar increase may represent

[47] One exception is Barr et al. (2022), who find that the effects of childhood exposure to the EITC on early-adulthood labor market outcomes are largely driven by men. This difference may reflect the timing of outcomes considered, as Barr et al. (2022) do not examine effects in the early twenties, which is precisely the period where I find the strongest impacts on women's labor market outcomes.

[48] As noted in Bitler and Figinski (2024), Food Stamps represent a large share of household food expenditure, which is broadly comparable to the 50–70% increase in household income generated by *PANES/AFAM-PE*.

different relative income changes, outcomes may be measured at different ages, and the focus on formal earnings in this paper implies that part of the estimated gains may reflect informal/formal income shifting.

The direction of the fertility effects is also consistent with recent evidence from income support policies in other contexts, including the EITC in the United States (Michelmore and Lopoo, 2021), conditional cash transfers in Latin America (e.g., Araujo and Macours, 2021; Attanasio et al., 2021; Barham et al., 2018), and cash transfer programs in Africa (Baird et al., 2011). The magnitude is broadly in line with the effects in Araujo and Macours (2021), who show that 1.5 additional years of exposure to PROGRESA increase age at first birth by 0.53 years, which is comparable to the roughly 0.5-year increase reported in Table 4 for a discontinuity of about two years of exposure. By contrast, effects appear larger than those documented for the United States: Michelmore and Lopoo (2021) find that an additional USD 1,000 in EITC exposure reduces the probability of giving birth by age 20 by 2–3%, compared to an estimated reduction of 11.6% per additional USD 1,000 before age 19 in this paper. Despite these differences in magnitude, a common qualitative pattern emerges across studies: the effects operate primarily through changes in fertility timing rather than shifts in fertility preferences along the extensive margin, as reflected in the U-shaped profile of effects over age.

A final set of results worth benchmarking concerns the effects on future welfare participation. Despite differences in settings, the null or negative effects on welfare applications, participation, and payments documented in this paper are closer to the findings in Deshpande (2016), who report no causal effect of SSI removal on intergenerational SSI receipt. This contrasts with evidence from Dahl et al. (2014); Dahl and Gielen (2021); Hartley et al. (2022), who document strong intergenerational transmission when studying disability insurance and AFDC/TANF in the United States. One possible explanation is that program design matters. In particular, conditionalities may be more effective in promoting human capital investments, thereby weakening some of the structural mechanisms underlying intergenerational welfare dependence. At the same time, more research is needed to better understand the role of conditionalities and other program features in shaping human capital accumulation and intergenerational outcomes.

5.4 Marginal Value of Public Funds

To assess the welfare effects of the program, I next calculate the Marginal Value of Public Funds (MVPF) associated with *PANES/AFAM-PE* as in Hendren and Sprung-Keyser (2020). While the analysis in this paper follows individuals up to 2023, behavioral responses to childhood income support may continue to accrue later in life. Hence, I first estimate the MVPF using outcomes observed up to 2023, which relies on fewer assumptions and are based exclusively on causal effects identified in the data. Second, I project MVPF estimates up to retirement age under alternative assumptions about the evolution of formal labor earnings over the life cycle. To maintain consistency with the empirical analysis, I present the main MVPF estimates for the average household in the *main sample*, defined as households within 2 p.p. of the eligibility threshold. The average household is comprised of 4.3 members, of whom 2.4 are children, with 48.6% being girls (see Appendix C). In addition, I will take 28 years, i.e., the average age of the sample when they are last observed, as a representative age in the analysis.

Willingness to Pay for *PANES/AFAM-PE*. First, I calculate beneficiaries' willingness to pay (WTP) for an additional dollar of program expenditure. A household-level WTP should aggregate the valuation of both adults and children who benefited from the program. Since *PANES/AFAM-PE* consists of a cash transfer with few restrictions on how it can be spent, it is reasonable to assume that adults value the cash transfer at face value. For children, I assume their WTP is given by the time-discounted increase in after-tax earnings induced by the transfer. In particular, I take the cumulative effect reported in column (4) of Table 2 (1.740), apply an average tax rate of 0.32, as in Bergolo and Cruces (2021), and a 3% time discount rate.[49] Throughout, I impose that the program has strictly zero earnings effects for men, and aggregate these individual-level gains to the household level by scaling them by the average number of girls per household at application, equal to 1.176 ($2.4 \times 49\%$). This leads to a household-level WTP for an additional dollar of *PANES/AFAM-PE* of $2.07, comprised of $1 for adults and $1.07 for children.

Fiscal Externalities. Fiscal externalities associated with *PANES/AFAM-PE* should capture changes in government revenues and expenditures induced by behavioral responses to the program. For simplicity, the baseline scenario focuses on first-order effects on payroll tax revenues arising from changes in formal labor earnings. On the adult side, as shown in Figure 1 and documented in Bergolo and Cruces (2021), the income eligibility thresholds used in *PANES/AFAM-PE* may have induced behavioral responses in parental labor supply. To account for these effects, I replicate the main estimation strategy based on Equation (1) using parental labor earnings as the outcome.[50] The estimates indicate that an additional USD 1,000 of income support during children's childhood is associated with a reduction of USD 383 in parents' cumulative earnings per parent over the analysis period.[51] I translate estimated cumulative effects on parents' formal earnings into fiscal externalities by applying the same payroll tax rate (0.32), aggregating across adults within the household (1.9), and discounting at a 3% rate. For children, increased payroll tax revenues from formal earnings are computed following the same procedure as for their WTP.

Static MVPF estimates. Based on these calculations, the estimated MVPF is 3.1. This implies that, up to the age 28 (or 18 years after the program started) each additional dollar of net government spending on *PANES/AFAM-PE* has provided nearly $3.1 of benefits for beneficiary households.

At this point, it is important to acknowledge that this estimate relies on several assumptions. The baseline scenario treats all estimated effects as real earnings responses, rather than shifts from informal to formal labor income, and abstracts from spillovers to other tax bases, such as VAT. To assess sensitivity to these assumptions, I simulate alternative scenarios that vary the share of real responses among children and adults and allow for additional fiscal externalities

[49] Full results, including details on the assumptions and the calculation of the MVPF, are reported in Appendix H.

[50] For consistency with the child-level analysis, I use the same endogenous variable when estimating effects on parents. In practice, I measure cumulative earnings for household head and partner in the *main sample*, and estimate child–parent level regressions weighted by the inverse of the number of children in the household. This weighting prevents individuals from households with more children from receiving disproportionate weights.

[51] While this cumulative effect is not statistically significant (p-value = 0.362), estimates are negative and statistically significant at shorter horizons, such as cumulative earnings measured after 5 or 10 years, suggesting that parental labor supply responses are concentrated early on and fade out as children age and households exit the program.

through VAT revenues.[52] The conservative scenario assumes that adults' responses are real labor-supply responses, while children's responses reflect only formalization. The optimistic scenario assumes the opposite: children's responses are real, while adults' responses reflect only formalization. I also consider an intermediate scenario in which 50% of the responses are real for both groups. Across all scenarios, children's WTP is generated only by real earnings gains, so purely shifting responses generate fiscal externalities but no WTP. Under these alternative assumptions, MVPF estimates range from 1.3 in the conservative scenario to 4.2 in the optimistic one. Thus, in all cases, the benefits accrued by beneficiary households up to age 28 exceed the monetary cost incurred by the government.

These estimates are conservative for various reasons. First, they abstract from several potential responses that could lead to reductions in government expenditures, such as lower *PANES/AFAM-PE* payments due to reduced future welfare participation and reduced health care costs associated with lower teen fertility.[53] Second, they abstract from potential changes in life expectancy, which have been shown to be quantitatively important for MVPF calculations (Bailey et al., 2024). Given that *PANES/AFAM-PE* substantially improves women's labor market outcomes and living standards, it is plausible that the program may also have increased life expectancy among female beneficiaries. Third, and perhaps more importantly, the estimates are based only on outcomes observed up to age 28 and therefore ignore potential effects accruing later in life. Next, I extend the analysis by forecasting earnings and revenues beyond age 28.

Dynamic MVPF estimates. To account for longer-run effects, I follow the long-term earnings projection approach in Hendren and Sprung-Keyser (2020). More specifically, I project *PANES/AFAM-PE* ineligible women's income beyond age 28 using the cross-sectional age profile of formal labor earnings observed in the 2023 National Household Survey (ECH), rescaled to match the average income level of ineligible women near the eligibility cutoff, and assuming a 0.5% annual wage growth rate. Eligible women are then assumed to follow the same lifecycle earnings profile, but at a proportionally higher level, with the earnings gap pinned down by the 2SLS estimate for the latest observed labor market outcome reported in column (2) of Table 2 (5.2%). This projection imposes a constant proportional earnings gap over the lifecycle, which may overstate long-run gains if effects decline over time. To assess the sensitivity of the results to this assumption, I also consider scenarios where the gap decays at annual rates ranging from 0.5% to 10%. These are conservative scenarios, as they allow only for gradual *declines* in earnings differences and do not account for potentially faster and steeper post-birth earnings recovery documented in the paper, driven by shifts in child-penalty trajectories.

Using these projected earnings paths, I compute the discounted stream of additional payroll tax revenues generated up to age 67 and accumulate fiscal externalities over time, including those already observed at age 28 in the baseline scenario. Under no decay, the additional payroll tax revenues generated by female beneficiaries' formal earnings fully offset the initial fiscal cost of an additional USD 1 transferred to the household around age 37. By retirement age, the program generates fiscal externalities of USD 2.4 per dollar transferred, implying net revenue

[52] For this, I assume a consumption share of 80% and apply the standard 22% VAT rate.

[53] I also abstract from potential changes in education expenditures. Given the null/small-scale effects on secondary and tertiary education enrollment, these effects are unlikely to be quantitatively important.

gains of USD 1.4 per dollar transferred. This also holds, when allowing the earnings gap to decay over time, with cumulative fiscal externalities remaining large even under relatively fast decay rates such as slightly below 10% yearly. Overall, forecasting dynamic earnings responses beyond 2023 substantially strengthens the case for investing in *PANES/AFAM-PE*.[54]

6 Conclusion

While gender gaps in labor market outcomes have narrowed over recent decades, large differences persist, and motherhood remains a central driver of these gaps. At the same time, a growing literature shows that income support during childhood can have lasting effects on adult outcomes, especially for women, but the mechanisms through which such policies shape women's labor market trajectories remain unclear. In this paper, I use novel and exhaustive administrative records from Uruguay's *PANES/AFAM-PE* program to estimate the causal effects of additional income support during childhood on labor market outcomes, as well as on intermediate outcomes related to education and fertility decisions. By following individuals across multiple stages of the life cycle and tracing how these outcomes evolve over time, I document that additional resources during childhood lead women to adopt more career-oriented transitions, characterized by higher educational attainment, earlier and more sustained labor-market attachment, and delayed entry into motherhood, which ultimately translate into higher cumulative earnings and improved living standards.

These results have direct implications for the evaluation and design of social safety net policies. First, the fact that additional income support during childhood generates economically large labor-market gains for women, relative to the size of the transfers, suggests that childhood income support can be a powerful and cost-effective policy instrument for improving women's economic outcomes. These effects are especially relevant for gender inequality, given that motherhood penalties account for a sizable share of gender differences in earnings. By influencing the timing of early-life events such as childbirth, income support policies can have lasting consequences for labor-market gender gaps, even when they are not explicitly designed with this objective in mind. This may be particularly important in contexts where motherhood penalties are larger for low-income mothers, such as Uruguay (Querejeta and Bucheli, 2022). In this sense, childhood income support may also promote social mobility by improving the long-run economic trajectories of women from disadvantaged households.

Furthermore, when choosing between alternative policy instruments, income support programs may offer advantages relative to policies that target specific margins, such as fertility (e.g., contraceptive access) or education (e.g., schooling subsidies), more directly. By relaxing household budget constraints, additional income during childhood can affect a broader set of outcomes, including children's health, nutrition, and education (Fiszbein et al., 2009; Bastagli et al., 2016; Almond et al., 2018; Molina Millán et al., 2019; Page, 2024), parental investments (Gennetian et al., 2024; Krause et al., 2025), and potentially generate positive spillovers (Egger

[54] One caveat with this dynamic approach is that the analysis abstracts from potential pension-related fiscal externalities. Higher formal employment and earnings could increase future public pension expenditures, while lower formal employment, weaker pension coverage, and lower lifetime income could lead to increased participation in welfare programs targeted to the elderly, higher public health care costs due to worse health at older ages, and potentially earlier mortality.

et al., 2022), while preserving flexibility in how resources are allocated and avoiding some of the distortions associated with narrowly targeted or in-kind interventions. Importantly, while cash transfer programs may be costly in the short run, the MVPF analysis shows that *PANES/AFAM-PE* pays for itself before age 40 and generates additional net fiscal gains thereafter.

Finally, several questions remain open. As the cohorts studied in this paper age, future work can assess whether the labor-market gains documented here persist further into adulthood and whether childhood income support affects completed fertility. At the same time, the heterogeneity in responses across cohorts highlights the importance of understanding how the timing of income support interact with fertility control, educational opportunities, and other institutional features. Addressing how early-life policies operate across different stages of the life cycle and across institutional environments therefore remains an important direction for future research.

Main Tables and Figures

Figure 1: First Stage and Validity of the RDD

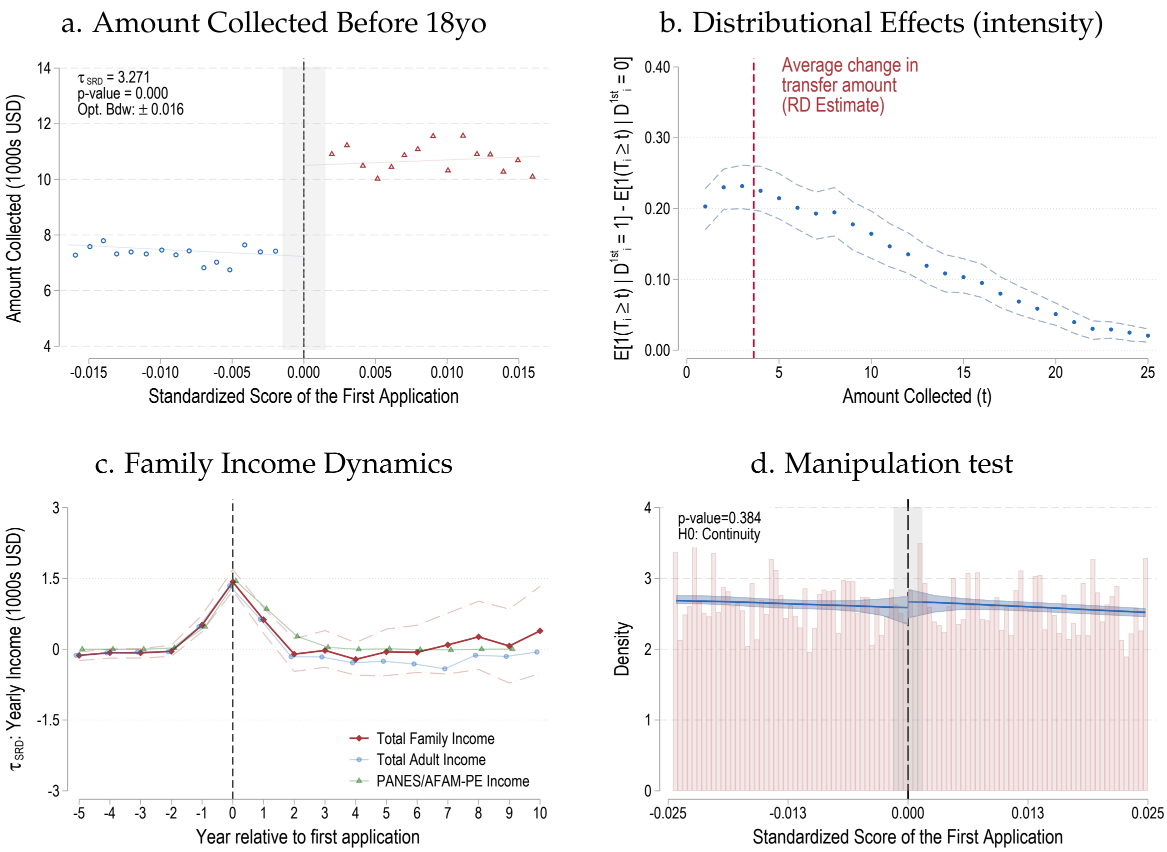

Notes: This figure reports a series of tests on the validity of the RDD design. Panel (a) shows the average total benefits received from *PANES/AFAM-PE* before age eighteen, expressed in 2008 USD (1000s), as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest; after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. Total benefits are computed as the cumulative value of program transfers received by households the individual belonged to prior to turning eighteen and are winsorized at the 95th percentile. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Average total benefits are grouped into 15 quantile-spaced bins on each side of the threshold (i.e., about three times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). In addition, the figure reports the point estimate of the local difference in average total benefits at the threshold (τ_{SRD}), based on Equation (2), together with the optimal bandwidth and the continuity test p-value, shown in the upper-left corner. Panel (b) illustrates that the first-stage estimates capture average differences in treatment intensity across the distribution, following Rose and Shem-Tov (2021). Each bin corresponds to a separate estimation of Equation (2), using the baseline technical specifications, where the outcome is defined as the indicator $1(T_i \geq t)$ for cumulative transfers received over the entire period. These estimates are obtained by varying t from 1,000 to 25,000 and estimating the corresponding difference in the probability of receiving at least t in transfers around the eligibility cutoff, i.e., local differences in $E[1(T_i \geq t) \mid D_i^{1st} = 1] - E[1(T_i \geq t) \mid D_i^{1st} = 0]$. Panel (c) reports the dynamic reduced-form estimates of household income components around the first program application. For each year relative to the first application, Equation (3) is estimated separately using a sharp RDD, following the baseline technical specification, where the outcome is alternatively defined as yearly total family income, yearly total adult income, or yearly transfers received from *PANES/AFAM-PE*. Total family income is defined as the sum of formal labor earnings of all baseline household members plus transfers, while total adult income includes formal labor earnings of baseline adult household members plus transfers. The figure plots the resulting local discontinuities at the eligibility threshold (τ_{SRD}), expressed in 2008 USD (1000s), winsorized at the 95th percentile. Panel (d) provides an illustration of a continuity test of Z_i^{1st} at the eligibility threshold, as proposed by Cattaneo et al. (2018). Parameter selection is based on the default options in the *rddensity* Stata command, and the p-value of the continuity test is reported in the upper-left corner. In all cases, the sample used corresponds to the *main sample* as defined in Section 3. In this case, I do not exclude individuals whose standardized poverty score at first application lies within 0.0015 to avoid creating holes in the distribution. Whenever shown, confidence intervals correspond to the 95% level, and inference is based on robust standard errors clustered at the household level.

Figure 2: Labor Market Outcomes - Women

Notes: This figure reports a series of reduced-form RDD estimates for *women* in the *main sample*, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Each panel illustrates the outcome of interest as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest; after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Outcomes are grouped into 15 quantile-spaced bins on each side of the threshold (i.e., about three times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). In addition, the figure reports the point estimate of the local difference in the outcome of interest at the threshold (τ_{SRD}), based on Equation (3), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. Panels (a) and (b) report effects on labor market outcomes measured in the last year observed, namely months worked and total labor earnings between August 2022 and July 2023. During this period, individuals in the sample are on average 28 years old; see Figure C.2 for additional descriptive statistics. Panels (c) and (d) report effects on cumulative outcomes, defined as the sum of months worked and total labor earnings accumulated over the entire observation period. Income/earnings variables, are expressed in 1000s of USD and are winsorized at the 95th percentile. Inference is based on robust standard errors clustered at the household level. Full estimates are reported in Table 2.

Figure 3: Robustness and Sensitivity Checks: Cumulative Earnings, Women

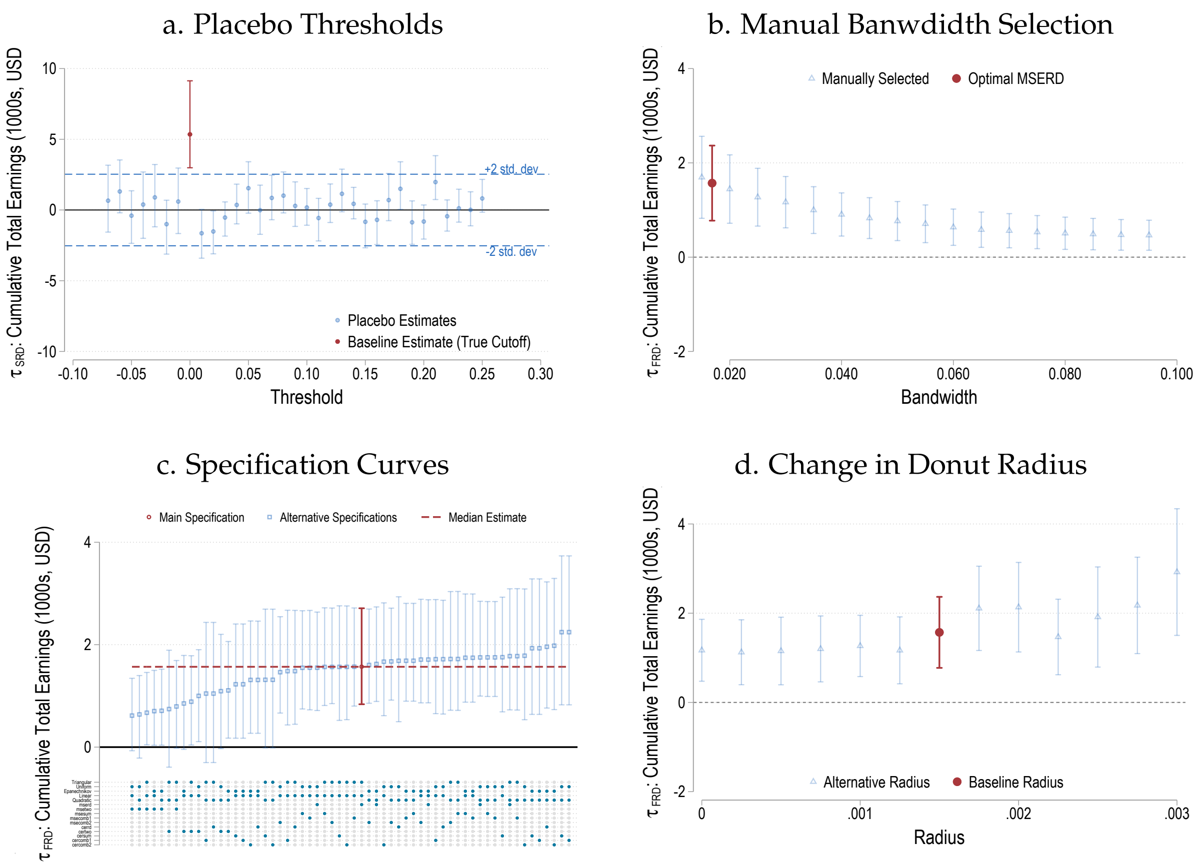

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *women's* cumulative earnings, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level. Earnings are expressed in 1000s of USD and are winsorized at the 95th percentile.

Figure 4: Labor Market Outcomes - Men

Notes: This figure reports a series of reduced-form RDD estimates for *men* in the *main sample* as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Each panel illustrates the outcome of interest as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest; after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Outcomes are grouped into 15 quantile-spaced bins on each side of the threshold (i.e., about three times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). In addition, the figure reports the point estimate of the local difference in the outcome of interest at the threshold (τ_{SRD}), based on Equation (3), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. Panels (a) and (b) report effects on labor market outcomes measured in the last year observed, namely months worked and total labor earnings between August 2022 and July 2023. During this period, individuals in the sample are on average 28 years old; see Figure C.2 for additional descriptive statistics. Panels (c) and (d) report effects on cumulative outcomes, defined as the sum of months worked and total labor earnings accumulated over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Inference is based on robust standard errors clustered at the household level. Full estimates are reported in Table 2.

Figure 5: Changes in Transition to Adulthood, Women

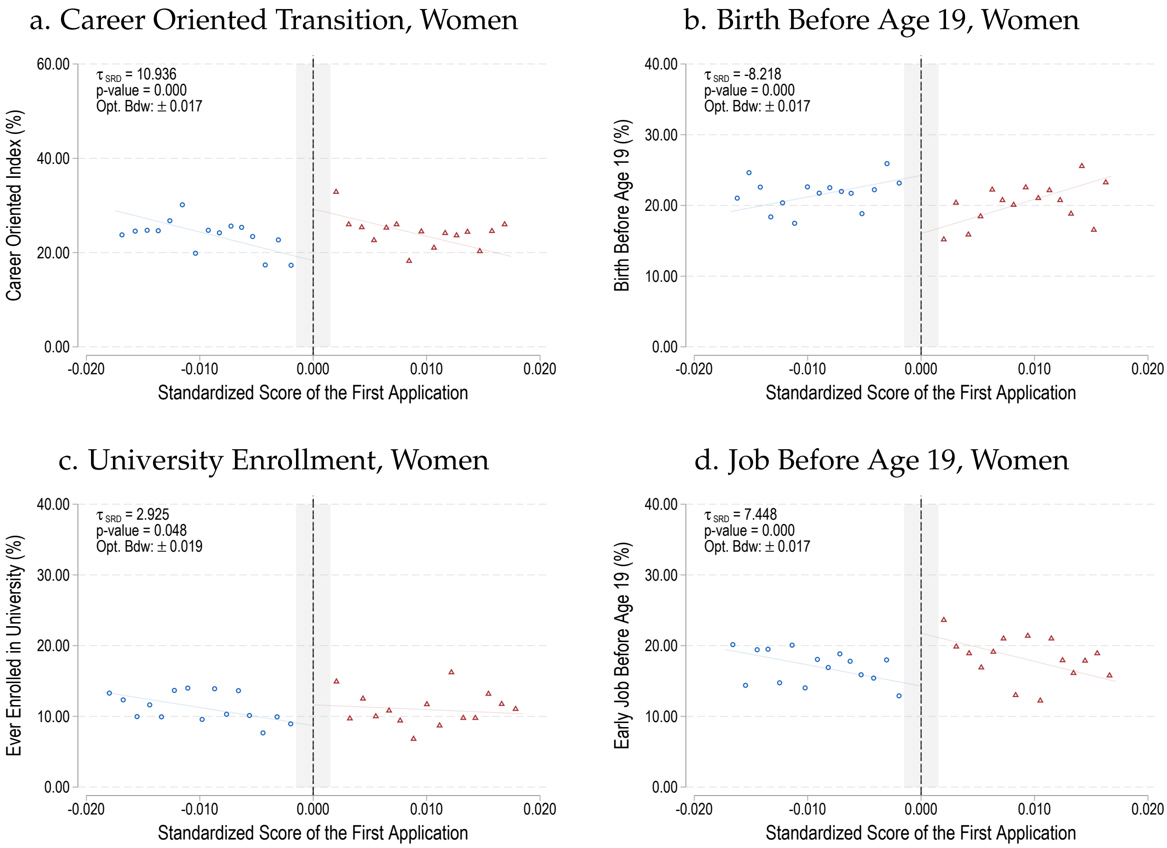

Notes: This figure reports a series of reduced-form RDD estimates for *women* in the *main sample* as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Each panel illustrates the outcome of interest as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest; after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Outcomes are grouped into 15 quantile-spaced bins on each side of the threshold (i.e., about three times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). In addition, the figure reports the point estimate of the local difference in the outcome of interest at the threshold (τ_{SRD}), based on Equation (3), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. Each panel illustrates one of the four main outcome variables discussed in Section 4.2. Panel (a) corresponds to an index of career-oriented transition to adulthood. This index is defined as a binary variable (0, 100) that equals 100 if the individual follows a career-oriented transition to adulthood, and 0 otherwise. An individual is classified as career-oriented if two conditions hold: (i) they did not have a child before age 19, and (ii) they either held their first relatively stable job before age 19, defined as an employment spell of at least four consecutive months, or enrolled in university. Panel (b) presents analogous analysis for a binary variable (0,100) that indicates having a child before age 19, panel (c) for the probability of ever being enrolled in university as a binary variable (0, 100), and panel (d) for a binary variable (0, 100) that indicates if the individual experienced a spell of at least four consecutive months of employment before age 19. Inference is based on robust standard errors clustered at the household level. Full estimates are reported in Table 3.

Figure 6: Dynamic Effects on Fertility and Employment Outcomes

a. Age by Age Effects, Women b. Age by Age Effects, Men

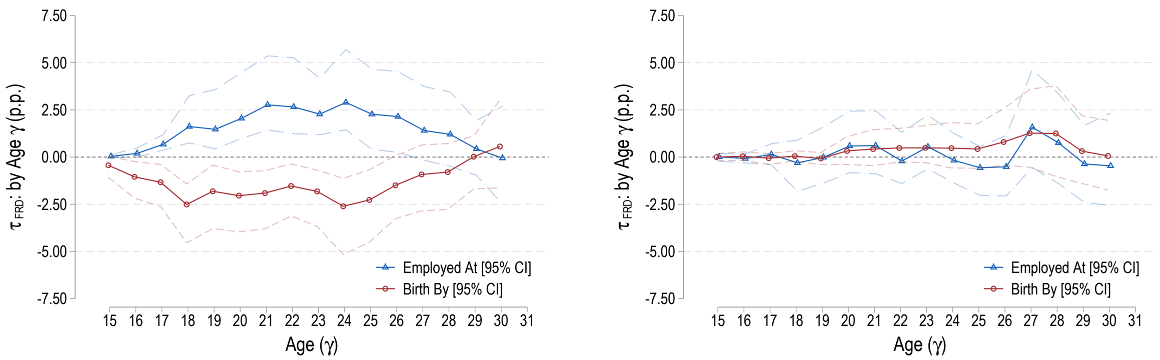

Notes: This figure reports dynamic 2SLS estimates of employment and fertility outcomes for women and men, respectively, as a function of age. The vertical axis reports the 2SLS estimate of the outcome measured at age γ, while the horizontal axis denotes age γ over the range 15 to 30. Each estimate τ_{FRD} is obtained from the baseline 2SLS specification in Equation (1) and can be interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Estimates are based on the *main sample*, excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold, and restricting the sample to individuals who are at least γ years old when last observed. For ages $\gamma < 18$, the endogenous variable is defined as the total amount of program transfers collected prior to age γ. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Employment outcomes are defined as a binary variable (0, 100) that equals 100 if the individual is employed at age γ, measured as having at least four consecutive months of employment at that age, and 0 otherwise. Fertility outcomes are defined as a binary variable (0, 100) that equals 100 if the individual has given birth by age γ, and 0 otherwise. Confidence intervals correspond to the 95 percent level, and inference is based on robust standard errors clustered at the household level. Full estimates are reported in Appendix E.

Figure 7: Child Penalty by Age of First Birth

a. Income Dynamics by Age of First Birth b. Income Dynamics by Age of First Birth, Pooled

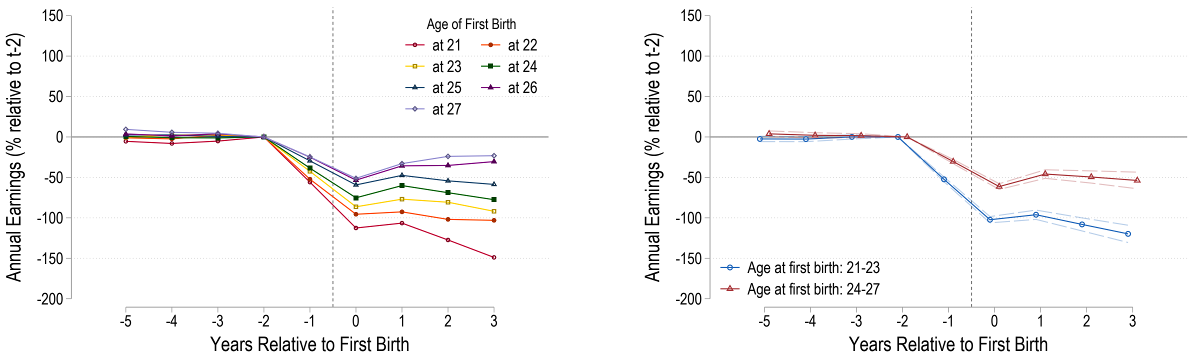

Notes: This Figure reports stacked difference-in-differences estimates of earnings dynamics around childbirth for women as explained in Section 4.3. Estimates are based on Equation (4), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Each panel shows estimated coefficients $\beta_{s,l}$ as a function of years relative to first birth (l), where the horizontal axis denotes event time and the vertical axis reports the estimated effect on earnings. In both panels, the outcome of interest is annual labor earnings, expressed as a percentage of the within-stack average earnings two years before birth ($t = -2$), which is the omitted reference period. Panel (a) reports separate estimates by age at first birth for women who give birth between ages 21 and 27. Panel (b) reports analogous estimates for aggregated age-at-first-birth groups, pooling ages 21-23 and 24-27 to increase statistical precision. Estimates are obtained using weights constructed to recover a weighted average treatment effect across age-at-first-birth sub-events, with weights proportional to cohort sample shares (Wing et al., 2024; Sun and Abraham, 2021). Whenever shown, confidence intervals correspond to the 95% level, and inference is based on robust standard errors clustered at the individual level. Further technical details on the stacked DiD construction, as well as full estimates, are provided in Appendix **??**.

Figure 8: Dynamics of Earnings Around Year of First Birth

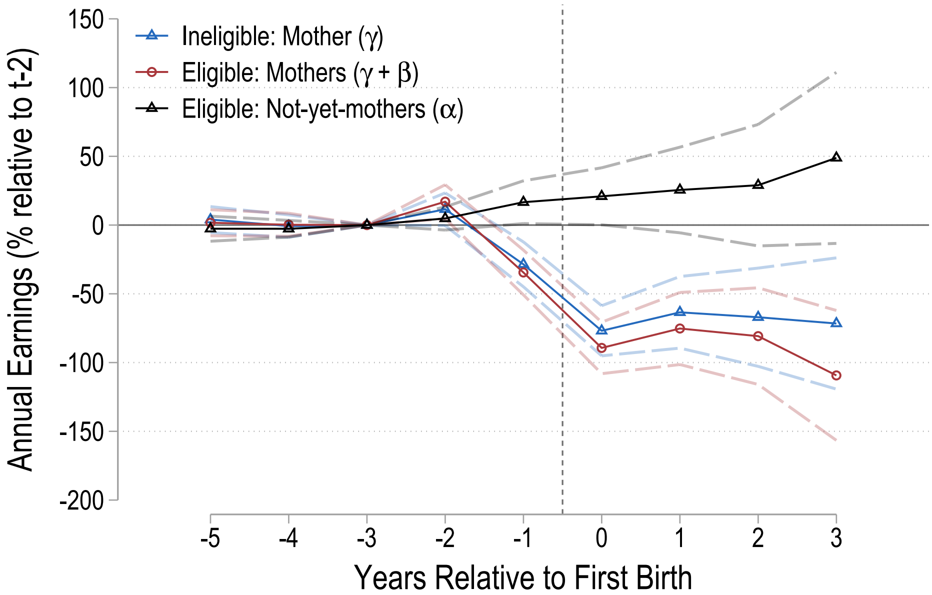

Notes: This Figure reports stacked difference-in-differences RDD estimates of earnings dynamics around childbirth for women, as explained in Section 4.3. Estimates are based on Equation (5). The sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Because of the RDD component of the specification, the sample is further restricted to observations within 5 p.p. of the eligibility threshold and excludes individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). The figure displays three series of coefficients as a function of years relative to first birth (l): the child penalty for ineligible mothers (γ_l), the child penalty for eligible mothers ($\gamma_l + \beta_l$), and the eligibility effect among not-yet-mothers (α_l). The horizontal axis denotes event time relative to first birth, and the vertical axis reports the estimated effect on earnings. The outcome of interest is annual labor earnings, expressed as a percentage of the within-stack average earnings two years before birth ($l = -2$), which is the omitted reference period. Estimates are obtained using weights constructed to recover a weighted average treatment effect across age-at-first-birth sub-events, with weights proportional to cohort sample shares (Wing et al., 2024; Sun and Abraham, 2021). Whenever shown, confidence intervals correspond to the 95% level, and inference is based on robust standard errors clustered at the individual level. Full estimates are provided in Appendix F.

Figure 9: Heterogeneous Effects by Age of Exposure

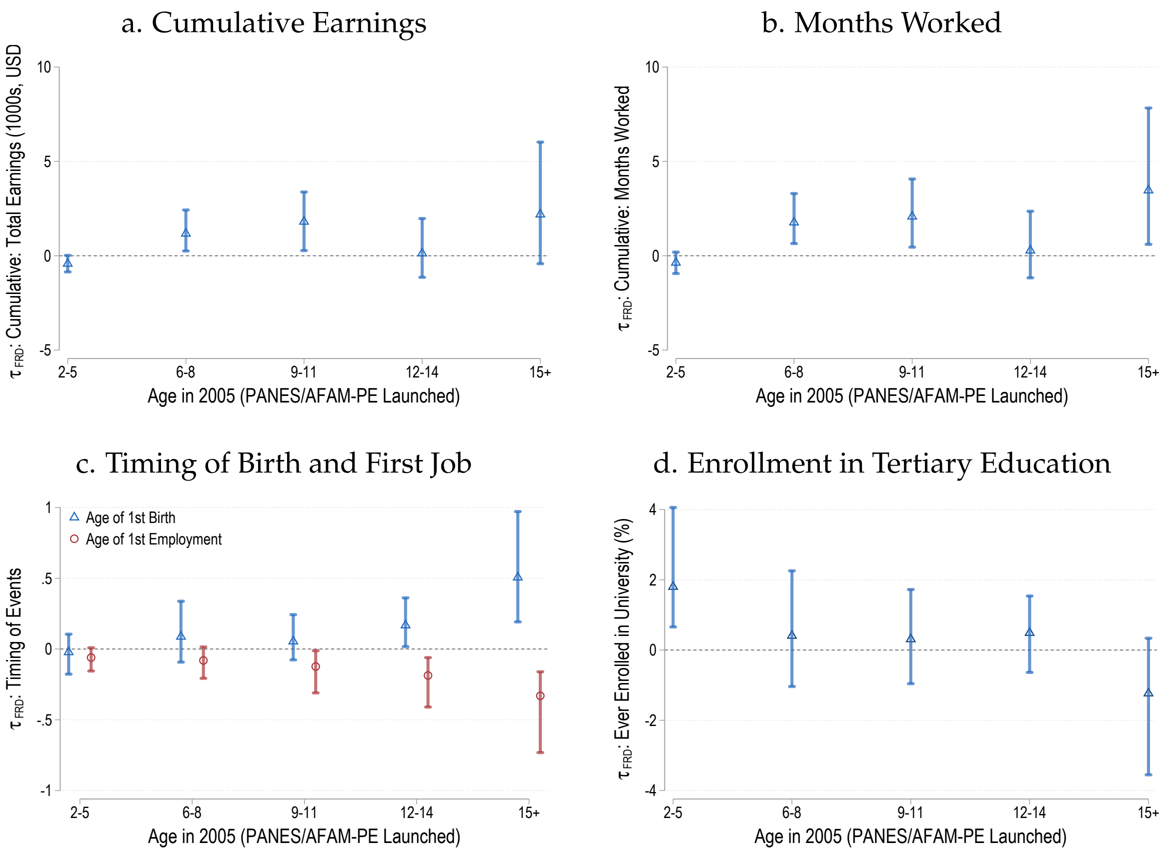

Notes: This Figure reports 2SLS estimates on a series of outcomes for women in the *main sample*, as defined in Section 3, excluding those with standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Each panel reports estimates separately for five cohorts defined by age at the rollout of *PANES/AFAM-PE*: 2-5, 6-8, 9-11, 12-14, and 15+, represented in the horizontal axis. The vertical axis reports the fuzzy RDD treatment effect (τ_{FRD}) based on Equation (1). Panels (a) and (b) report effects on cumulative outcomes, defined as the sum of months worked and total labor earnings accumulated over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Panel (c) reports analogous estimates for age-at-transition outcomes: age of first birth, and age of first employment spell. Panel (d) reports estimates for university enrollment, defined as a binary variable (0, 100) that equals 100 if the individual ever enrolled in university, and 0 otherwise. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level. Full estimates are reported in Appendix F.

Figure 10: Effects by Having Early Birth and Education Level

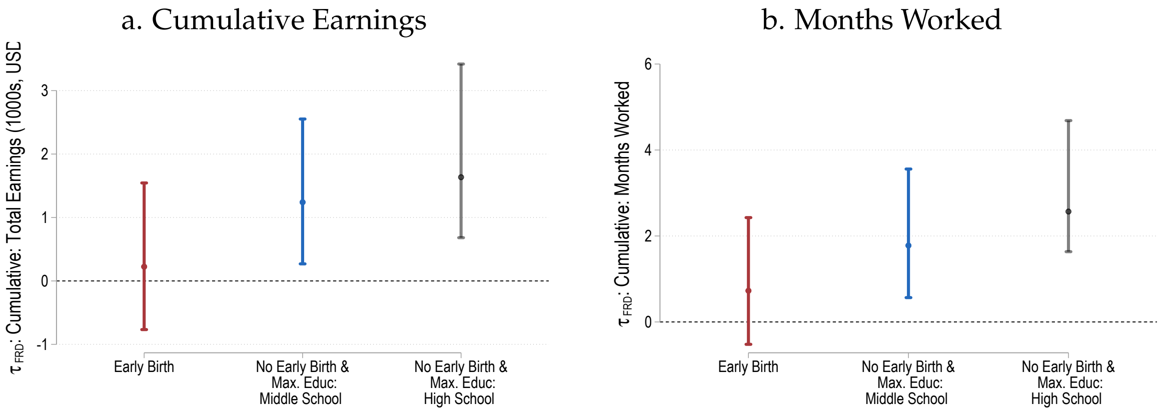

Notes: This Figure reports 2SLS estimates on cumulative labor market outcomes outcomes for women in the *main sample*, as defined in Section 3, excluding those with standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Each panel reports estimates separately for three groups represented on the horizontal axis: (i) individuals who had a child before age 19, (ii) individuals without an early birth who did not enroll beyond lower secondary education, and (iii) individuals without an early birth who enrolled in upper secondary or higher education. The vertical axis reports the fuzzy RDD treatment effect (τ_{FRD}) based on Equation (1). Panels (a) and (b) report effects on cumulative earnings and months worked, respectively, defined as the sum accumulated over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Confidence intervals correspond to the 95% level, and inference is based on robust standard errors clustered at the household level. Full estimates are reported in Appendix F.

Table 1: Continuity at Baseline

	Ineligible Adj. Mean	Eligible Intercept	Difference (2) - (1)	p-value Robust
	(1)	(2)	(3)	(4)
Predicted Eligibility	0.66	0.66	-0.002	0.740
HH - Number of people	3.98	3.94	-0.047	0.551
HH - Single Parent	57.37	54.28	-3.091	0.121
HH - Mother's Age when First born	22.52	22.56	0.041	0.778
HH - Avg. Age	19.39	21.00	1.614	0.000
HH - Baseline App. Before Sep.2005	5.63	7.22	1.589	0.756
HHH - Employed	63.12	64.38	1.264	0.373
HHH - Years of Educ.	7.18	7.03	-0.148	0.118
HHH - Self-reported Labor Income	0.15	0.15	-0.001	0.759
Female	49.70	48.59	-1.113	0.195
Number of applications	2.71	2.79	0.082	0.102
Age at 1st. Application	4.67	4.67	0.000	0.797
Age at 31 Dec. 2021	19.50	19.50	-0.003	0.761

Notes: This table reports estimates on the continuity of baseline covariates $\mathbf{X}$ at the eligibility threshold for individuals in the *main sample*, as defined in Section 3, excluding those with standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Estimates are based on the reduced form specification included in Equation (3) including all fixed effects $\mathbf{\Lambda}$ but no additional covariates $\mathbf{X}$ as these are the outcomes of interest. Each row represents a different dependent variable. Importantly, row (1) presents a summary test where the outcome, i.e., predicted eligibility, is the predicted probability of being just above the eligibility threshold based on all the other baseline characteristics included in the table. Column (1) depicts the intercept for the local linear regression fitted in the ineligible side of the centered threshold (i.e., $z < 0$), while column (2) does the same for the local linear regression fitted in the eligible side of the threshold. Column (3) is simply the difference between columns (1) and (2). Columns (4) reports the robust p-values for the continuity tests. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Inference is based on robust standard errors clustered at the household level.

Table 2: Effects on Labor Market Outcomes

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.148*	0.199**	2.520***	1.740***	0.161**	0.142	0.607	0.252
	(0.084)	(0.088)	(0.636)	(0.551)	(0.073)	(0.085)	(0.558)	(0.494)
Robust p-value	0.079	0.029	0.000	0.002	0.039	0.115	0.237	0.561
Mean Baseline Outcome	4.95	3.85	37.33	24.74	5.24	4.69	47.91	34.84
Effect Size (%)	2.98%	5.16%	6.75%	7.03%	3.07%	3.04%	1.27%	0.72%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.485*	0.654**	8.311***	5.738***	0.521**	0.464	1.968	0.817
	(0.270)	(0.282)	(1.978)	(1.753)	(0.232)	(0.273)	(1.796)	(1.597)
Robust p-value	0.074	0.025	0.000	0.002	0.033	0.107	0.228	0.555
Mean Baseline Outcome	4.95	3.85	37.33	24.74	5.24	4.69	47.91	34.84
Effect Size (%)	9.80%	16.98%	22.27%	23.19%	9.95%	9.89%	4.11%	2.34%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.284***	3.293***	3.299***	3.298***	3.240***	3.257***	3.240***	3.242***
	(0.275)	(0.278)	(0.285)	(0.290)	(0.230)	(0.238)	(0.229)	(0.230)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.50	7.52	7.55	7.54	7.35	7.34	7.36	7.35
Effect Size (%)	43.78%	43.79%	43.69%	43.72%	44.07%	44.40%	44.04%	44.09%
Selection of Bandwidth:								
Opt. Bandwidth	[0.017]	[0.017]	[0.016]	[0.016]	[0.019]	[0.018]	[0.019]	[0.019]
Effective Obs.	8,068	7,952	7,687	7,506	9,718	9,196	9,729	9,653

Notes: This table reports estimates of the effects of additional childhood income support on labor market outcomes. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3), corresponding to the sharp discontinuities shown in Figures 2 through 4. Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 3: Summary of Effects on Transitions to Adulthood

	Women				Men			
	Birth Before Age 19 (1)	University Enrollment (2)	Job Before Age 19 (3)	Career-oriented Index (4)	Birth Before Age 19 (5)	University Enrollment (6)	Job Before Age 19 (7)	Career-oriented Index (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	-2.511***	0.910*	2.285***	3.363***	0.051	-0.090	0.393	-0.022
	(0.643)	(0.482)	(0.639)	(0.725)	(0.116)	(0.236)	(0.452)	(0.608)
Robust p-value	0.000	0.053	0.000	0.000	0.675	0.679	0.410	0.702
Mean Baseline Outcome	21.68	11.25	17.07	23.68	2.05	4.90	32.04	34.42
Effect Size (%)	-11.58%	8.09%	13.39%	14.21%	2.47%	-1.84%	1.23%	-0.06%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	-8.218***	2.925**	7.448***	10.936***	0.161	-0.291	1.254	-0.071
	(1.915)	(1.522)	(1.966)	(2.130)	(0.367)	(0.761)	(1.441)	(1.959)
Robust p-value	0.000	0.048	0.000	0.000	0.669	0.677	0.395	0.702
Mean Baseline Outcome	21.68	11.25	17.07	23.68	2.05	4.90	32.04	34.42
Effect Size (%)	-37.90%	26.01%	43.63%	46.19%	7.84%	-5.93%	3.91%	-0.21%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.273***	3.215***	3.259***	3.252***	3.169***	3.224***	3.194***	3.223***
	(0.276)	(0.257)	(0.271)	(0.268)	(0.131)	(0.181)	(0.160)	(0.213)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.41	7.35	7.37	7.38	6.84	7.16	7.04	7.20
Effect Size (%)	44.14%	43.74%	44.23%	44.06%	46.36%	45.02%	45.38%	44.78%
Selection of Bandwidth:								
Opt. Bandwidth	[0.017]	[0.019]	[0.017]	[0.017]	[0.048]	[0.028]	[0.034]	[0.022]
Effective Obs.	8,139	8,974	8,320	8,447	26,267	14,957	18,247	11,481

Notes: This table reports estimates of the effects of additional childhood income support on transition-to-adulthood outcomes. Columns (1) through (4) correspond to women. Column (1) reports effects on the probability of having children before age 19 as a binary variable (0, 100), column (2) on the probability of ever being enrolled in university as a binary variable (0, 100), and column (3) reports effects on the probability of having a spell of at least four consecutive months of employment before age 19 as a binary variable (0, 100). Column (4) reports estimates for the career-oriented index. This index is defined as a binary variable (0, 100) that equals 100 if the individual follows a career-oriented transition to adulthood, and 0 otherwise. An individual is classified as career-oriented if two conditions hold: (i) they did not have a child before age 19, and (ii) they either held their first relatively stable job before age 19, defined as an employment spell of at least four consecutive months, or enrolled in university. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 4: Effects on Age of Events and Extensive Margin

	Women				Men			
	Ever Had a Child (1)	Age of 1st. Birth (2)	Ever Employed (3)	Age of 1st. Job (4)	Ever Had a Child (5)	Age of 1st. Birth (6)	Ever Employed (7)	Age of 1st. Job (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	-1.039	0.160***	-0.372	-0.263***	-0.149	-0.254***	0.169	-0.018
	(0.646)	(0.049)	(0.678)	(0.054)	(0.527)	(0.089)	(0.562)	(0.051)
Robust p-value	0.135	0.001	0.667	0.000	0.822	0.008	0.726	0.780
Mean Baseline Outcome	55.14	20.00	71.80	20.74	23.70	22.88	78.00	19.61
Effect Size (%)	-1.88%	0.80%	-0.52%	-1.27%	-0.63%	-1.11%	0.22%	-0.09%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	-3.322	0.502***	-1.217	-0.847***	-0.480	-0.755***	0.546	-0.053
	(2.030)	(0.149)	(2.219)	(0.159)	(1.697)	(0.255)	(1.818)	(0.152)
Robust p-value	0.129	0.001	0.666	0.000	0.821	0.006	0.722	0.780
Mean Baseline Outcome	55.14	20.00	71.80	20.74	23.70	22.88	78.00	19.61
Effect Size (%)	-6.02%	2.51%	-1.69%	-4.08%	-2.02%	-3.30%	0.70%	-0.27%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.196***	3.141***	3.271***	3.219***	3.222***	2.970***	3.237***	2.997***
	(0.249)	(0.199)	(0.271)	(0.286)	(0.218)	(0.341)	(0.228)	(0.252)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.41	6.34	7.47	6.34	7.20	5.17	7.37	6.53
Effect Size (%)	43.14%	49.51%	43.77%	50.77%	44.77%	57.48%	43.91%	45.88%
Selection of Bandwidth:								
Opt. Bandwidth	[0.019]	[0.031]	[0.017]	[0.019]	[0.021]	[0.022]	[0.019]	[0.018]
Effective Obs.	9,401	8,593	8,210	6,278	11,099	2,780	9,864	6,938

Notes: This table reports estimates of the effects of additional childhood income support on extensive margin and age-of-event-type variables. Columns (1) through (4) correspond to women. Column (1) reports effects on the probability of ever having children (0, 100), column (2) on the age at first birth, column (3) reports effects on the probability of ever having a spell of at least four consecutive months of employment, and column (4) reports estimates for age of the first employment spell. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Age of first birth and age of first job estimates are restricted to women and men who had children and ever had an employment spell, respectively. Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 5: Effects on Labor Market Mobility and Stability - Women

	Mobility History					Stability in 2023m7		
	Number of Firms (1)	Stayed Same Firm (2)	Low Mobility (3)	Moderate Mobility (4)	High Mobility (5)	Tenure in current job (6)	Log. Tenure in current job (7)	Tenure +2 Year in current job (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.003	-1.295**	-1.632**	1.798**	-0.161	0.858	0.034**	2.606***
	(0.034)	(0.578)	(0.872)	(0.855)	(0.819)	(0.577)	(0.017)	(1.052)
Robust p-value	0.859	0.029	0.036	0.027	0.900	0.117	0.024	0.007
Mean Baseline Outcome	2.69	19.75	38.61	28.39	33.24	38.81	3.21	44.16
Effect Size (%)	0.11%	-6.56%	-4.23%	6.33%	-0.48%	2.21%		5.90%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.009	-4.116**	-5.267**	5.948**	-0.519	2.679	0.107**	7.933***
	(0.109)	(1.820)	(2.764)	(2.787)	(2.640)	(1.796)	(0.052)	(3.112)
Robust p-value	0.859	0.027	0.033	0.025	0.901	0.108	0.021	0.006
Mean Baseline Outcome	2.69	19.75	38.61	28.39	33.24	38.81	3.21	44.16
Effect Size (%)	0.35%	-20.84%	-13.64%	20.95%	-1.56%	6.90%		17.96%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.181***	3.177***	3.228***	3.309***	3.223***	3.122***	3.113***	3.045***
	(0.231)	(0.223)	(0.292)	(0.312)	(0.291)	(0.226)	(0.245)	(0.306)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.53	6.30	6.23	6.33	6.23	5.89	5.94	6.41
Effect Size (%)	42.23%	50.42%	51.77%	52.30%	51.73%	52.97%	52.44%	47.52%
Selection of Bandwidth:								
Opt. Bandwidth	[0.022]	[0.026]	[0.018]	[0.017]	[0.018]	[0.041]	[0.035]	[0.024]
Effective Obs.	10,520	9,098	6,276	5,733	6,310	6,983	6,053	4,835

Notes: This table reports estimates of the effects of additional childhood income support on labor market mobility and job stability outcomes for women. Columns (1) through (5) report measures of labor market mobility over the full observation period. Column (1) reports the number of different firms in which the individual held her main position, with never employed coded as 0. Column (2) reports the probability of remaining in the same firm when last observed, conditional on ever being employed. Columns (3), (4), and (5) report binary indicators of low, moderate, and high mobility, respectively, conditional on ever being employed. Low mobility corresponds to having one or two main employers over the entire period, moderate mobility to having three or four employers, and high mobility to having five or more. Columns (6) through (8) report measures of job stability conditional on being employed in July 2023. Column (6) reports tenure in the current job, column (7) reports log tenure in the current job, and column (8) reports a binary indicator equal to 100 if tenure in the current job is at least two years, which corresponds approximately to the median tenure in the data. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 6: Effects on Labor Market Outcomes by Mother vs Non-Mothers

	Non-Mothers				Mothers			
	University Enrollment (1)	Age of 1st Job (2)	Cumulative: Months Employed (3)	Cumulative: Earnings (4)	University Enrollment (5)	Age of 1st. Job (6)	Cumulative: Months Employed (7)	Cumulative: Earnings (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	1.585*	-0.236***	2.749***	1.860***	-0.125	-0.226***	1.856***	1.176**
	(0.882)	(0.071)	(0.844)	(0.708)	(0.271)	(0.055)	(0.634)	(0.511)
Robust p-value	0.055	0.000	0.000	0.006	0.494	0.000	0.004	0.022
Mean Baseline Outcome	19.96	20.29	34.70	24.92	4.09	21.01	39.64	25.07
Effect Size (%)	7.94%	-1.16%	7.92%	7.47%	-3.06%	-1.07%	4.68%	4.69%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	5.131*	-0.752***	8.900***	5.931***	-0.386	-0.720***	5.899***	3.702**
	(2.750)	(0.196)	(2.469)	(2.135)	(0.836)	(0.169)	(1.974)	(1.589)
Robust p-value	0.051	0.000	0.000	0.004	0.487	0.000	0.003	0.017
Mean Baseline Outcome	19.96	20.29	34.70	24.92	4.09	21.01	39.64	25.07
Effect Size (%)	25.70%	-3.70%	25.65%	23.81%	-9.44%	-3.43%	14.88%	14.77%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.237***	3.182***	3.238***	3.188***	3.085***	3.189***	3.178***	3.149***
	(0.416)	(0.461)	(0.439)	(0.421)	(0.177)	(0.216)	(0.208)	(0.202)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	8.50	7.74	8.67	8.72	6.29	5.25	6.41	6.37
Effect Size (%)	38.07%	41.13%	37.33%	36.55%	49.03%	60.76%	49.60%	49.46%
Selection of Bandwidth:								
Opt. Bandwidth	[0.020]	[0.021]	[0.018]	[0.019]	[0.038]	[0.032]	[0.029]	[0.030]
Effective Obs.	4,336	3,040	3,849	4,106	10,393	6,270	7,962	8,410

Notes: This table reports estimates of the effects of additional childhood income support on education and labor market outcomes by motherhood status. Columns (1) through (4) correspond to women who have not had children by the end of the observation period (non-mothers), while columns (5) through (8) correspond to women who are mothers. Column (1) reports the probability of ever being enrolled in university, defined as a binary variable (0, 100). Column (2) reports age at first employment and is conditional on ever being employed. Columns (3) and (4) report cumulative labor market outcomes over the entire observation period, namely cumulative months employed and cumulative labor earnings, respectively. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous outcomes for mothers. Estimates are conducted separately for non-mothers and mothers, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 7: Summary of Effects on Secondary Education Outcomes

	Women				Men			
	Ever Enrolled (1)	Years Enrolled (2)	Enrolled in 12th grade (3)	Max. Grade Enrolled (4)	Ever Enrolled (5)	Years Enrolled (6)	Enrolled in 12th grade (7)	Max. Grade Enrolled (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	-0.053	0.006	2.041***	0.070**	0.294	0.045	0.065	0.047*
	(0.563)	(0.048)	(0.786)	(0.030)	(0.683)	(0.037)	(0.483)	(0.026)
Robust p-value	0.843	0.998	0.008	0.025	0.844	0.334	0.838	0.083
Mean Baseline Outcome	80.02	3.99	28.99	10.18	64.31	2.62	11.49	9.34
Effect Size (%)	-0.07%	0.15%	7.04%		0.46%	1.72%	0.56%	
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	-0.169	0.019	6.686***	0.237**	0.952	0.145	0.210	0.151*
	(1.802)	(0.157)	(2.494)	(0.096)	(2.216)	(0.119)	(1.560)	(0.082)
Robust p-value	0.843	0.998	0.007	0.021	0.840	0.329	0.839	0.078
Mean Baseline Outcome	80.02	3.99	28.99	10.18	64.31	2.62	11.49	9.34
Effect Size (%)	-0.21%	0.49%	23.07%		1.48%	5.53%	1.82%	
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.202***	3.277***	3.276***	3.372***	3.243***	3.219***	3.229***	3.231***
	(0.252)	(0.279)	(0.287)	(0.306)	(0.243)	(0.221)	(0.238)	(0.249)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.39	7.43	7.42	7.68	7.15	7.23	7.16	7.82
Effect Size (%)	43.35%	44.09%	44.15%	43.88%	45.34%	44.50%	45.12%	41.33%
Selection of Bandwidth:								
Opt. Bandwidth	[0.019]	[0.017]	[0.016]	[0.017]	[0.018]	[0.021]	[0.019]	[0.025]
Effective Obs.	9,227	8,014	7,678	6,700	9,307	10,830	9,632	8,441

Notes: This table reports estimates of the effects of additional childhood income support on secondary education outcomes. Columns (1) through (4) correspond to women. Column (1) reports the probability of ever being enrolled in secondary education, defined as a binary variable (0, 100). Column (2) reports the total number of years enrolled in secondary education. Column (3) reports the probability of ever enrolling in 12th grade, defined as a binary variable (0, 100). Column (4) reports the maximum grade attained, from 6 to 12. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 8: Mediation Analysis

	Dep. Var.: Cumulative: Total Earnings (1000s, USD)							
	All Women				Women Who Gave Birth			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	1.740***	1.349***	1.528***	1.309***	1.176**	0.692	0.957**	0.645
	(0.551)	(0.498)	(0.518)	(0.486)	(0.511)	(0.456)	(0.457)	(0.429)
Robust p-value	0.002	0.009	0.004	0.009	0.022	0.115	0.030	0.108
Mean Baseline Outcome	24.74	24.80	24.70	24.73	25.07	25.11	25.11	25.13
Effect Size (%)	7.03%	5.44%	6.19%	5.30%	4.69%	2.76%	3.81%	2.57%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	5.738***	4.526***	5.068***	4.416***	3.702**	2.166	2.977**	2.010*
	(1.753)	(1.641)	(1.677)	(1.614)	(1.589)	(1.420)	(1.412)	(1.334)
Robust p-value	0.002	0.007	0.003	0.007	0.017	0.103	0.025	0.098
Mean Baseline Outcome	24.74	24.80	24.70	24.73	25.07	25.11	25.11	25.13
Effect Size (%)	23.19%	18.25%	20.52%	17.86%	14.77%	8.62%	11.85%	8.00%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.298***	3.355***	3.317***	3.372***	3.149***	3.129***	3.111***	3.116***
	(0.290)	(0.274)	(0.290)	(0.280)	(0.202)	(0.188)	(0.187)	(0.181)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.54	7.50	7.55	7.54	6.37	6.33	6.33	6.32
Effect Size (%)	43.72%	44.73%	43.96%	44.70%	49.46%	49.40%	49.16%	49.30%
Selection of Bandwidth:								
Opt. Bandwidth	[0.016]	[0.017]	[0.016]	[0.016]	[0.030]	[0.034]	[0.034]	[0.036]
Effective Obs.	7,506	8,053	7,488	7,843	8,410	9,473	9,517	10,004
Mediators:								
Early Birth	No	Yes	No	Yes	No	No	No	No
Age at 1st Birth	No	No	No	No	No	Yes	No	Yes
High Educ.	No	No	Yes	Yes	No	No	Yes	Yes

Notes: This table reports estimates of the effects of additional childhood income support on cumulative labor earnings. The outcome of interest is cumulative labor earnings accumulated over the entire observation period, expressed in thousands of USD and winsorized at the 95th percentile. Columns (1) through (4) correspond to all women, while columns (5) through (8) correspond to women who gave birth during the observation period. Column (1) reports estimates from the baseline specification and is identical to column (4) in Table 2. Columns (2) through (4) sequentially augment the baseline specification by including mediator variables. Column (2) includes an indicator for early childbirth, defined as having the first birth at or before age 19. Column (3) includes an indicator for high educational attainment, defined as ever being enrolled in upper secondary education (grades 10, 11, or 12). Column (4) includes both mediators simultaneously. Columns (5) through (8) follow an analogous structure for mothers. Column (5) reports the baseline specification for mothers and is identical to column (8) in Table 6. Column (6) includes age at first birth as a mediator. Column (7) includes the indicator for high educational attainment, and column (8) includes both mediators jointly. Estimates are conducted separately for all women and women who gave birth, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table 9: Welfare Participation as Adults

	Women					Men				
	Ever Applied (1)	Age 1st. Application (2)	Ever Accepted (3)	Number of Payments (4)	Poverty Score (5)	Ever Applied (6)	Age 1st. Application (7)	Ever Accepted (8)	Number of Payments (9)	Poverty Score (10)
a. Fuzzy RDD Estimate (τ_{FRD})										
Amount Collected (1000s, USD)	-0.713	0.137*	-0.818	-1.268*	-0.011***	0.222	-0.073	0.066	0.130	-0.002
	(0.765)	(0.064)	(0.740)	(0.571)	(0.004)	(0.503)	(0.081)	(0.478)	(0.343)	(0.004)
Robust p-value	0.458	0.074	0.352	0.057	0.006	0.422	0.381	0.532	0.518	0.664
Mean Baseline Outcome	35.58	21.82	32.78	22.03	0.23	17.50	23.35	16.22	9.36	0.29
Effect Size (%)	-2.00%	0.63%	-2.49%	-5.76%	-4.80%	1.27%	-0.31%	0.41%	1.39%	-0.73%
b. Sharp RDD Estimate (τ_{SRD})										
Elig. 1st. App.	-2.334	0.422*	-2.677	-4.153*	-0.034***	0.724	-0.233	0.214	0.421	-0.007
	(2.498)	(0.193)	(2.414)	(1.838)	(0.011)	(1.640)	(0.259)	(1.550)	(1.113)	(0.013)
Robust p-value	0.459	0.060	0.352	0.054	0.003	0.420	0.371	0.531	0.515	0.647
Mean Baseline Outcome	35.58	21.82	32.78	22.03	0.23	17.50	23.35	16.22	9.36	0.29
Effect Size (%)	-6.56%	1.94%	-8.17%	-18.85%	-14.94%	4.14%	-1.00%	1.32%	4.50%	-2.35%
c. First Stage (τ_{SRD})										
Elig. 1st. App.	3.275***	3.077***	3.275***	3.274***	3.114***	3.262***	3.209***	3.244***	3.248***	3.212***
	(0.294)	(0.237)	(0.294)	(0.292)	(0.244)	(0.248)	(0.297)	(0.243)	(0.245)	(0.298)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.43	6.64	7.43	7.43	6.66	7.18	5.47	7.15	7.15	5.48
Effect Size (%)	44.10%	46.33%	44.09%	44.09%	46.74%	45.46%	58.68%	45.35%	45.40%	58.63%
Selection of Bandwidth:										
Opt. Bandwidth	[0.015]	[0.036]	[0.015]	[0.016]	[0.034]	[0.017]	[0.037]	[0.018]	[0.018]	[0.037]
Effective Obs.	7,439	6,366	7,437	7,497	6,018	9,014	3,394	9,305	9,219	3,364

Notes: This table reports estimates of the effects of additional childhood income support on welfare participation outcomes in adulthood. Importantly, an individual is identified as applying as an adult if they submit an application form under a different household identifier and are listed as the household head or as the spouse or partner of the household head. Columns (1) through (5) correspond to women. Column (1) reports the probability of applying for welfare benefits (*PANES/AFAM-PE*) as an adult, defined as a binary variable (0, 100). Column (2) reports age at first application as an adult and is conditional on having applied. Column (3) reports the probability of being accepted as an adult beneficiary, defined as a binary variable (0, 100). Column (4) reports the number of payments received as an adult. Column (5) reports the adult poverty score at application and is conditional on having applied. Columns (6) through (10) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

References

Adda, J., C. Dustmann, and K. Stevens (2017). The career costs of children. *Journal of Political Economy 125*(2), 293–337.

Aizer, A., S. Eli, J. Ferrie, and A. Lleras-Muney (2016, April). The Long-Run Impact of Cash Transfers to Poor Families. *American Economic Review 106*(4), 935–971.

Aizer, A., H. Hoynes, and A. Lleras-Muney (2022, May). Children and the US Social Safety Net: Balancing Disincentives for Adults and Benefits for Children. *Journal of Economic Perspectives 36*(2), 149–174.

Akee, R. K. Q., W. E. Copeland, G. Keeler, A. Angold, and E. J. Costello (2010, January). Parents' Incomes and Children's Outcomes: A Quasi-experiment Using Transfer Payments from Casino Profits. *American Economic Journal: Applied Economics 2*(1), 86–115.

Almond, D., J. Currie, and V. Duque (2018, December). Childhood Circumstances and Adult Outcomes: Act II. *Journal of Economic Literature 56*(4), 1360–1446.

Amarante, V., M. Ferrando, and A. Vigorito (2013). Teenage school attendance and cash transfers: An impact evaluation of panes. *Economía 14*(1), 61–96.

Amarante, V., M. Manacorda, E. Miguel, and A. Vigorito (2016, May). Do Cash Transfers Improve Birth Outcomes? Evidence from Matched Vital Statistics, and Program and Social Security Data. *American Economic Journal: Economic Policy 8*(2), 1–43.

Angelov, N., P. Johansson, and E. Lindahl (2016). Parenthood and the gender gap in pay. *Journal of labor economics 34*(3), 545–579.

Angrist, J. D. and G. W. Imbens (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association 90*(430), 431–442.

Araujo, M. C. and K. Macours (2021). Education, Income and Mobility: Experimental Impacts of Childhood Exposure to Progresa after 20 Years. Technical Report IDB-WP-01288, IDB.

Arnett, J. J. (2000). Emerging adulthood: A theory of development from the late teens through the twenties. *American Psychologist 55*, 469–480. Place: US Publisher: American Psychological Association.

Attanasio, O., L. C. Sosa, C. Medina, C. Meghir, and C. M. Posso-Suárez (2021). Long term effects of cash transfer programs in Colombia. Technical Report w29056, National Bureau of Economic Research.

Bailey, D., G. J. Duncan, C. L. Odgers, and W. Yu (2017). Persistence and fadeout in the impacts of child and adolescent interventions. *Journal of research on educational effectiveness 10*(1), 7–39.

Bailey, M. J. (2006). More power to the pill: The impact of contraceptive freedom on women's life cycle labor supply. *The quarterly journal of economics 121*(1), 289–320.

Bailey, M. J., H. Hoynes, M. Rossin-Slater, and R. Walker (2024). Is the social safety net a long-term investment? large-scale evidence from the food stamps program. *Review of Economic Studies 91*(3), 1291–1330.

Bailey, M. J., V. W. Lang, A. Prettyman, I. Vrioni, L. J. Bart, D. Eisenberg, P. Fomby, J. Barber, and V. Dalton (2023). How costs limit contraceptive use among low-income women in the us: a randomized control trial. Technical report, National Bureau of Economic Research.

Baird, S., C. McIntosh, and B. Ozler (2011, November). Cash or Condition? Evidence from a Cash Transfer Experiment. *The Quarterly Journal of Economics 126*(4), 1709–1753.

Barham, T., K. Macours, and J. A. Maluccio (2018). Experimental Evidence of Exposure to a Conditional Cash Transfer During Early Teenage Years: Young Women's Fertility and Labor Market Outcomes. Technical report, CEPR Discussion Papers.

Barr, A., J. Eggleston, and A. A. Smith (2022). Investing in Infants: the Lasting Effects of Cash Transfers to New Families*. *The Quarterly Journal of Economics Forthcoming*, qjac023.

Bastagli, F., J. Hagen-Zanker, L. Harman, V. Barca, G. Sturge, T. Schmidt, and L. Pellerano (2016). Cash Transfers: What Does the Evidence Say? A Rigourus Review of Programme Impact and the Role of Design and Implementation Features. Tech. Rep., Overseas Dev. Inst., London, Overseas Development Institute.

Bastian, J., L. Bian, and J. Grogger (2022). How Did Safety-Net Reform Affect the Education of Adolescents from Low-Income Families? *Labour Economics 77*, 102031.

Bastian, J. and K. Michelmore (2018, October). The Long-Term Impact of the Earned Income Tax Credit on Children's Education and Employment Outcomes. *Journal of Labor Economics 36*(4), 1127–1163. Publisher: The University of Chicago Press.

Bergolo, M. and G. Cruces (2021). The anatomy of behavioral responses to social assistance when informal employment is high. *Journal of Public Economics 193*, 104313.

Bergolo, M. and E. Galván (2018, March). Intra-household Behavioral Responses to Cash Transfer Programs. Evidence from a Regression Discontinuity Design. *World Development 103*, 100–118.

Berthelon, M. E. and D. I. Kruger (2011, February). Risky behavior among youth: Incapacitation effects of school on adolescent motherhood and crime in Chile. *Journal of Public Economics 95*(1), 41–53.

Bertrand, M., C. Goldin, and L. F. Katz (2010). Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American economic journal: applied economics 2*(3), 228–255.

Bitler, M. and T. F. Figinski (2024). Long-run effects of food assistance: Evidence from the food stamp program and administrative data. Technical report, National Bureau of Economic Research.

Black, S. E., P. J. Devereux, and K. G. Salvanes (2008). Staying in the Classroom and Out of the Maternity Ward? The Effect of Compulsory Schooling Laws on Teenage Births. *The Economic Journal 118*(530), 1025–1054.

Blau, F. D. and L. M. Kahn (2017, September). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature 55*(3), 789–865.

Bloomfield, J. and J. M. Cabrera (2026). Long-term impacts on education of a cash transfer during early life. *Journal of Policy Analysis and Management 45*(1), e22663.

Bulman, G., R. Fairlie, S. Goodman, and A. Isen (2021, April). Parental Resources and College Attendance: Evidence from Lottery Wins. *American Economic Review 111*(4), 1201–1240.

Cabella, W. and C. Velázquez (2022). Abortion Legalization in Uruguay: Effects on Adolescent Fertility. *Studies in Family Planning 53*(3), 491–514.

Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik (2019). Regression Discontinuity Designs Using Covariates. *The Review of Economics and Statistics 101*(3), 442–451.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica 82*(6), 2295–2326. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA11757.

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2019). *A practical introduction to regression discontinuity designs: Foundations*. Cambridge University Press.

Cattaneo, M. D., M. Jansson, and X. Ma (2018). Manipulation Testing Based on Density Discontinuity. *The Stata Journal 18*(1), 234–261. Publisher: SAGE Publications Sage CA: Los Angeles, CA.

Ceni, R., C. Parada, I. Perazzo, and E. Sena (2021). Birth Collapse and a Large-Scale Access Intervention with Subdermal Contraceptive Implants. *Studies in Family Planning 52*(3), 321–342.

Chetty, R., J. N. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan (2011). How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star *. *The Quarterly Journal of Economics 126*(4), 1593–1660.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The quarterly journal of economics 129*(4), 1553–1623.

Cunha, F. and J. Heckman (2007). The Technology of Skill Formation. *American Economic Review 97*(2), 31–47.

Cunha, F., J. J. Heckman, and S. M. Schennach (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica 78*(3), 883–931.

Dahl, G. B. and A. C. Gielen (2021, April). Intergenerational Spillovers in Disability Insurance. *American Economic Journal: Applied Economics 13*(2), 116–150.

Dahl, G. B., A. R. Kostøl, and M. Mogstad (2014, November). Family Welfare Cultures *. *The Quarterly Journal of Economics 129*(4), 1711–1752.

De Chaisemartin, C. and X. d'Haultfoeuille (2023). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. *The econometrics journal 26*(3), C1–C30.

DeCicca, P. and H. Krashinsky (2020). Does education reduce teen fertility? evidence from compulsory schooling laws. *Journal of Health Economics 69*, 102268.

Deshpande, M. (2016, November). Does Welfare Inhibit Success? The Long-Term Effects of Removing Low-Income Youth from the Disability Rolls. *American Economic Review 106*(11), 3300–3330.

Doepke, M., A. Hannusch, F. Kindermann, and M. Tertilt (2023). The economics of fertility: A new era. In *Handbook of the Economics of the Family*, Volume 1, pp. 151–254. Elsevier.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. Walker (2022). General equilibrium effects of cash transfers: experimental evidence from kenya. *Econometrica 90*(6), 2603–2643.

Field, E. M., M. McKelway, and A. Voena (2026). Gender norms and development. Working Paper 34832, National Bureau of Economic Research.

Fiszbein, A., N. R. Schady, F. H. G. Ferreira, M. Grosh, N. Keleher, P. Olinto, and E. Skoufias (2009). *Conditional Cash Transfers: Reducing Present and Future Poverty*. Washington D.C: World Bank.

Gallen, Y., J. S. Joensen, E. R. Johansen, and G. F. Veramendi (2024). The labor market returns to delaying pregnancy. *Available at SSRN 4554407*.

Galván, E., C. Parada, M. Querejeta, and S. Salvador (2024). Gender gaps and family leaves in latin america. *Review of Economics of the Household 22*(2), 387–414.

Gennetian, L. A., G. J. Duncan, N. A. Fox, S. Halpern-Meekin, K. Magnuson, K. G. Noble, and H. Yoshikawa (2024). Effects of a monthly unconditional cash transfer starting at birth on family investments among us families with low income. *Nature human behaviour 8*(8), 1514–1529.

Goldin, C. (2004). The long road to the fast track: Career and family. *The Annals of the American Academy of Political and Social Science 596*(1), 20–35.

Goldin, C. (2006). The quiet revolution that transformed women's employment, education, and family. *American economic review 96*(2), 1–21.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American economic review 104*(4), 1091–1119.

Goldin, C. (2021). Career and family: Women's century-long journey toward equity.

Goldin, C. and L. F. Katz (2002). The power of the pill: Oral contraceptives and women's career and marriage decisions. *Journal of political Economy 110*(4), 730–770.

Gustafsson, S. (2001, June). Optimal age at motherhood. Theoretical and empirical considerations on postponement of maternity in Europe. *Journal of Population Economics 14*(2), 225–247.

Hahn, J., P. Todd, and W. Van der Klaauw (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica 69*(1), 201–209.

Hartley, R. P., C. Lamarche, and J. P. Ziliak (2022, March). Welfare Reform and the Intergenerational Transmission of Dependence. *Journal of Political Economy 130*(3), 523–565.

Heckman, J. J. (2006). Skill formation and the economics of investing in disadvantaged children. *Science 312*(5782), 1900–1902.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly journal of economics 135*(3), 1209–1318.

Hotz, V. J., S. W. McElroy, and S. G. Sanders (2005). Teenage childbearing and its life cycle consequences: Exploiting a natural experiment. *Journal of Human Resources 40*(3), 683–715.

Hoynes, H., D. W. Schanzenbach, and D. Almond (2016). Long-Run Impacts of Childhood Access to the Safety Net. *American Economic Review 106*(4), 903–934.

Imbens, G. W. and T. Lemieux (2008). Regression Discontinuity Designs: A Guide to Practice. *Journal of Econometrics 142*(2), 615–635.

Jepsen, C., P. Mueser, and K. Troske (2016). Labor Market Returns to the GED Using Regression Discontinuity Analysis. *Journal of Political Economy 124*(3), 621–649.

Keane, M. P. and K. I. Wolpin (2010). The Role of Labor and Marriage Markets, Preference Heterogeneity and the Welfare System in the Life Cycle Decisions of Black, Hispanic and White Women. *International Economic Review 51*(3), 851–892.

Kearney, M. S. and P. B. Levine (2009, February). Subsidized Contraception, Fertility, and Sexual Behavior. *The Review of Economics and Statistics 91*(1), 137–151.

Kearney, M. S. and P. B. Levine (2012). Why is the teen birth rate in the united states so high and why does it matter? *Journal of Economic Perspectives 26*(2), 141–166.

Kearney, M. S. and P. B. Levine (2014). Teen births are falling: What's going on. Technical report, Washington, DC: The Brookings Institution.

Kearney, M. S. and P. B. Levine (2015). Investigating recent trends in the us teen birth rate. *Journal of Health Economics 41*, 15–29.

Klepinger, D., S. Lundberg, and R. Plotnick (1999). How Does Adolescent Fertility Affect the Human Capital and Wages of Young Women? *The Journal of Human Resources 34*(3), 421.

Kleven, H., C. Landais, and G. Leite-Mariante (2025). The child penalty atlas. *Review of Economic Studies 92*(5), 3174–3207.

Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics 11*(4), 181–209.

Korting, C., C. Lieberman, J. Matsudaira, Z. Pei, and Y. Shen (2023). Visual inference and graphical representation in regression discontinuity designs. *The Quarterly Journal of Economics 138*(3), 1977–2019.

Krause, P. K., E. Rhodes, S. Miller, A. W. Bartik, D. E. Broockman, and E. Vivalt (2025). The impact of unconditional cash transfers on parenting and children. Technical report, National Bureau of Economic Research.

List, J. A. (2020). Non est disputandum de generalizability? a glimpse into the external validity trial. Technical report, National Bureau of Economic Research.

Londoño-Vélez, J. and E. Saravia (2025). The impact of being denied a wanted abortion on women and their children. *The Quarterly Journal of Economics 140*(2), 1061–1110.

MacLeod, W. B. and M. Urquiola (2019). Is Education Consumption or Investment? Implications for School Competition. *Annual Review of Economics 11*(1), 563–589. _eprint: https://doi.org/10.1146/annurev-economics-080218-030402.

Manacorda, M., E. Miguel, and A. Vigorito (2011, July). Government Transfers and Political Support. *American Economic Journal: Applied Economics 3*(3), 1–28.

Manoli, D. and N. Turner (2018, May). Cash-on-Hand and College Enrollment: Evidence from Population Tax Data and the Earned Income Tax Credit. *American Economic Journal: Economic Policy 10*(2), 242–271.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics 142*(2), 698–714.

Melentyeva, V. and L. Riedel (2023). Child penalty estimation and mothers' age at first birth. Technical report, ECONtribute Discussion Paper.

Michelmore, K. and L. M. Lopoo (2021, December). The Effect of EITC Exposure in Childhood on Marriage and Early Childbearing. *Demography 58*(6), 2365–2394.

Miller, A. R. (2011, July). The Effects of Motherhood Timing on Career Path. *Journal of Population Economics 24*(3), 1071–1100.

Miller, S., L. R. Wherry, and D. G. Foster (2023). The economic consequences of being denied an abortion. *American Economic Journal: Economic Policy 15*(1), 394–437.

Molina Millán, T., T. Barham, K. Macours, J. A. Maluccio, and M. Stampini (2019). Long-Term Impacts of Conditional Cash Transfers: Review of the Evidence. *The World Bank Research Observer 34*(1), 119–159.

Nicolau, R. (2023). Welfare stigma after take-up: Evidence from public cash transfers in uruguay. *Serie Documentos de Trabajo; 14/23*.

Oreopoulos, P. (2011). Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes. *American Economic Journal: Economic Policy 3*(4), 148–171.

Page, M. E. (2024). New advances on an old question: Does money matter for children's outcomes? *Journal of economic literature 62*(3), 891–947.

Parada, C. (2023). Cash transfers and intra-household decision-making in uruguay. *Journal of Family and Economic Issues 44*(3), 757–775.

Parada, C. (2024). Fertility responses to cash transfers in uruguay. *World Development Perspectives 33*, 100574.

Parker, S. W. and T. Vogl (2023). Do conditional cash transfers improve economic outcomes in the next generation? evidence from mexico. *The Economic Journal 133*(655), 2775–2806.

Querejeta, M. and M. Bucheli (2022, October). The Effect of Childbirth on Women's Formal Labour Market Trajectories: Evidence from Uruguayan Administrative Data. *The Journal of Development Studies 0*(0), 1–15.

Ramírez Leira, L. (2021). Segregación escolar público-privado por nivel socioeconómico en Uruguay: Un análisis en base a microdescomposiciones. Working Paper 275, Documento de Trabajo.

Rose, E. K. and Y. Shem-Tov (2021). How does incarceration affect reoffending? estimating the dose-response function. *Journal of Political Economy 129*(12), 3302–3356.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics 235*(2), 2218–2244.

Schmidt, L., T. Sobotka, J. Bentzen, A. Nyboe Andersen, and on behalf of the ESHRE Reproduction and Society Task Force (2012, January). Demographic and Medical Consequence of the Postponement of Parenthood. *Human Reproduction Update 18*(1), 29–43.

Settersten Jr, R. A., F. F. Furstenberg, and R. G. Rumbaut (2008, September). *On the Frontier of Adulthood: Theory, Research, and Public Policy*. University of Chicago Press.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics 225*(2), 175–199.

Todd, P. E. and K. I. Wolpin (2006). Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility. *American Economic Review 96*(5), 1384–1417.

Todd, P. E. and K. I. Wolpin (2008). Ex Ante Evaluation of Social Programs. *Annales d'Économie et de Statistique ""*(91/92), 263–291.

UNDP (2023). 2023 gender social norms index (gsni). *UNDP (United Nations Development Programme)*.

Wasserman, M. (2023a). Do long work hours constrain women's fertility choice?". Technical report, Working paper February.

Wasserman, M. (2023b). Hours constraints, occupational choice, and gender: Evidence from medical residents. *The Review of Economic Studies 90*(3), 1535–1568.

Wing, C., S. M. Freedman, and A. Hollingsworth (2024). Stacked difference-in-differences. Technical report, National Bureau of Economic Research.

Online Appendix

Delaying Fertility, Advancing Careers: The Lasting Consequences of Growing Up with a Safety Net

Matías Giaccobasso

May 8, 2026

A Further Details on the Institutional Background

A.1 Uruguay's Background and Comparison to Other Countries

This section compares Uruguay's socioeconomic conditions with those from a group of Latin American and OECD countries in 2005, the year in which *PANES/AFAM-PE* was introduced.

Table A.1 shows that, relative to other Latin American countries, Uruguay faced relatively favorable macroeconomic and institutional indicators. GDP per capita in PPP terms was $12,296 in Uruguay, comparable to Chile ($12,550) and Argentina ($13,465), and higher than in Brazil ($11,081), Colombia ($8,432), and Costa Rica ($9,899). In terms of human development, Uruguay's HDI was 0.778 (rank 53), ranking below Chile (0.804, rank 46) and Argentina (0.812, rank 44), but above most other Latin American countries, including Brazil (0.718, rank 80), Colombia (0.711, rank 84), Costa Rica (0.748, rank 65), and Mexico (0.736, rank 72). Uruguay also exhibited relatively high tax revenues as a share of GDP (22.9%), ranking behind Brazil (32.7%) and close to Argentina (24.5%), but above Chile (20.8%), Colombia (18.3%), Costa Rica (21.8%), and Mexico (10.9%). This places Uruguay among the Latin American countries with relatively high state capacity and a strong public sector presence.

Despite these favorable macroeconomic and fiscal indicators, Uruguay lagged substantially behind OECD countries along dimensions closely related to human capital accumulation and demographic outcomes. Table A.1 reports that the lower secondary education completion rate in Uruguay was 59.3%, comparable to Argentina (56.4%) and higher than in Mexico (47.7%), Brazil (43.1%), Colombia (41.0%), and Costa Rica (39.9%), but far below rates observed in OECD countries such as Sweden (89.1%), the United States (93.7%), Italy (68.8%), and Spain (62.5%). A similar pattern emerges for adolescent fertility: Uruguay's adolescent fertility rate was 60.3 births per 1,000 women aged 15–19, lower than in Colombia (88.1) and Mexico (72.8), but higher than in Chile (50.8), and markedly above rates in Spain (11.5), Sweden (5.7), and the United States (39.5).

A.2 PANES/AFAM-PE: Design and Implementation

Uruguay has one of the oldest and most developed social protection systems in Latin America. Key components such as pensions, maternity leave, sickness and disability insurance, and unemployment benefits were introduced relatively early in the 20th century. Family allowances, for instance, were introduced in 1943. However, despite their early introduction, these benefits were, until very recently, largely restricted to formal, registered workers.

PANES was introduced in April 2005 by a new center-left government (*Frente Amplio*) in response to the economic crisis of the early 2000s. Its main goal was to extend family income support to informal, unregistered workers. In 2008, the program was expanded and rebranded as AFAM-PE. In this paper, I treat PANES and its AFAM-PE expansion as a single program, hereafter *PANES/AFAM-PE*, as they share the same core objectives and a very similar design.

The program was designed with two main goals. In the short run, it aimed to mitigate the sharp increase in poverty caused by the economic crisis of 2001–2002. In 2005, the overall poverty rate in urban areas was 36.6%. Child poverty was even more severe, reaching 61.3% among children aged 0–5, 58.0% among children aged 6–12, and 49.3% among children aged 13–17. In the longer run, the program aimed to improve conditions for human capital accumulation and social mobility by conditioning cash transfers on education and health requirements. In this sense, the transfers were intended to help households move out of poverty sooner rather than later, while at the same time encouraging investments in children's human capital.

The fiscal cost of the program is estimated at around 0.4% of GDP, which is comparable to large-scale programs such as PROGRESA-Oportunidades in Mexico and Bolsa Família in Brazil, and not far from public spending on family benefits in the United States (0.67% of GDP).[55] The program was widely publicized, rapidly became the most generous anti-poverty program in the country's history, and remains in place today. I next describe its main features.

PANES: Plan de Asistencia Nacional a la Emergencia Social

The main component of *PANES* was a cash transfer targeted at the poorest 150,000 households in the country. The base amount was USD 133, expressed in January 2008 PPP terms.[56] In addition, the program provided a supplementary transfer between USD 29 and USD 78 to households with underage children, which were about 70% of the participant households. The base transfer alone corresponds to about 63.9% of self-reported household income at the time of application. When supplementary payments are added, this share increases to roughly 78.1%. These figures should nonetheless be interpreted with caution, as households may have had incentives to underreport income to the survey enumerators. As an alternative benchmark, in 2005 the monthly transfer received by a household with four or more children was equivalent to one minimum wage.

While the most important component of PANES was the cash transfer, the program also included several smaller components, such as temporary public employment programs, education and job training, and other minor interventions, including access to public utilities, building materials, and free dental and eye health care. Although 96.7% of participating households received the cash transfer, fewer than 20% took part in the remaining components. Therefore, *PANES* can be interpreted mainly as a cash transfer program, even though for a small share of households it may have provided a broader set of benefits.[57]

A total of 187,727 households applied to *PANES*, corresponding to 679,077 individuals. This represents about 17% of all households in Uruguay and roughly 20% of the total population. Eligibility for the program was determined using two criteria. First, applicant households were required to have a per-capita income below USD 131, a condition that affected only around 10% of applications. Household per-capita income was computed as the sum of individual incomes divided by the number of household members. Individual income was defined as the maximum between self-reported income at the time of the household interview and income registered with the Social Security Agency.

The second eligibility criterion was based on a proxy-means test. After applying to the program, each household was assigned a poverty score z. This score was constructed using information collected during an in-person household interview conducted by program officials.

[55] https://data.oecd.org/socialexp/family-benefits-public-spending.htm

[56] In local currency, this corresponded to UYU 1,360. In what follows, all income variables are converted to 2008 PPP using CPI and PPP conversion factors.

[57] See Manacorda et al. (2011) for more details on the program.

It represents the predicted probability that the household's per-capita income falls below a critical income level. The variables used to compute the score were:

- Indicator for public-sector employees in the household
- Indicator for pensioners
- Average years of education among individuals aged 18 and above
- Number of household members
- Indicator for children aged 0–5, 6–12, and 13–17
- Indicator for private health insurance coverage
- Residential overcrowding
- Toilet facilities
- Wealth index based on household durables

Households with a poverty score above a region-specific threshold were deemed eligible for the program, while those below the threshold were considered ineligible.[58] The proxy-means testing methodology and the region-specific eligibility thresholds were designed by researchers at the country's largest public university (*Universidad de la Republica*). Importantly, neither the income threshold nor the poverty score threshold, nor the variables used to calculate it, were disclosed to applicants, and households were not informed of their own poverty score.

Once accepted, participant households were required to satisfy school attendance, regular health check-ups, and per-capita income requirements. However, these conditions were not enforced in practice due to administrative constraints and lack of enforcement capacity. Moreover, there is no evidence that participants were removed from the program due to non-compliance with the established requirements. Further details on the model used to calculate the poverty score, as well as additional information on program implementation, can be found in Manacorda et al. (2011); Amarante et al. (2016).

AFAM-PE: Asignaciones Familiares - Plan de Equidad

On January 1, 2008, *PANES* was expanded and rebranded as *AFAM-PE*. While *PANES* was conceived as a temporary program, *AFAM-PE* extended its benefits to a larger share of the population and was permanently established as a central component of the social safety net. The program's main components, namely eligibility criteria and the type of benefits and conditionalities, remained largely unchanged, and the transition between the two phases was straightforward. Most *PANES* participants were automatically enrolled in *AFAM-PE*. In addition, households rejected during the first phase were automatically enrolled in the second phase if they met the new, more lenient eligibility requirements (i.e., a lower eligibility threshold).

By December 2023, the total number of applications had reached 796,131, corresponding to about 1,599,556 different individuals. This represents a substantial increase relative to the population covered by *PANES*, both because the program expanded its coverage and because it remained in place for a longer period. There were three main differences between *PANES* and *AFAM-PE*. First, *AFAM-PE* introduced the presence of underage children in the household

[58] Regions are defined as follows: Montevideo (capital city); North (Artigas, Salto, and Rivera); Center-North (Paysandú, Río Negro, Tacuarembó, Durazno, Treinta y Tres, and Cerro Largo); Center-South (Soriano, Florida, Flores, Lavalleja, and Rocha); and South (Colonia, San José, Canelones, and Maldonado).

as an eligibility requirement, while *PANES* imposed no such restriction. Second, the eligibility threshold became more lenient. Third, the formula used to determine the transfer amount was modified. Under the new structure, households received a baseline payment of USD 57 per child aged 0–17, subject to an equivalence scale of 0.6. In addition to the base payment, households received an extra USD 24 per child aged 12-17 enrolled in secondary education, also subject to the same equivalence scale.

Conditionalities for remaining in the program were also not enforced during the first years of implementation. However, in June 2013 the government began to require households to comply with program requirements. For example, in June 2013 more than 30,000 children were removed from the program due to non-compliance with education enrollment requirements. In subsequent years, enforcement depended strongly on the administration in charge and was relatively intermittent.

Table A.1: Uruguay's Background and Comparison to Other Countries (2005)

Country	GDP PPP	HDI	HDI Rank	Tax Revenues (% gdp)	Completed Lower Secondary	Adolescent Fertility (per 1000s)	Unemployment rate
Argentina	13,465	0.812	44	24.50	56.38	63.55	11.51
Brazil	11,081	0.718	80	32.70	43.10	75.97	10.55
Chile	12,550	0.804	46	20.80	76.00	50.78	9.34
Colombia	8,432	0.711	84	18.30	40.99	88.12	11.06
Costa Rica	9,899	0.748	65	21.80	39.90	67.14	6.75
Italy	30,138	0.876	23	39.10	68.76	6.77	7.73
Mexico	13,260	0.736	72	10.90	47.68	72.77	3.56
OECD (Avg.)	29,854			32.60		32.68	6.83
Spain	27,630	0.857	29	35.20	62.52	11.50	9.15
Sweden	34,198	0.911	8	47.40	89.14	5.66	7.81
United States	44,123	0.907	13	26.10	93.70	39.47	5.08
Uruguay	12,296	0.778	53	22.90	59.26	60.26	12.01

Notes: GDP per capita (PPP, current international dollars) is from the World Bank's World Development Indicators (International Comparison Program). Human Development Index (HDI) values and HDI ranks are from the United Nations Development Programme (UNDP), Human Development Report 2025. Tax revenues (percent of GDP) are from OECD Global Revenue Statistics. Lower secondary completion rates (population aged 25 and over) are from the UNESCO Institute for Statistics. Adolescent fertility rates (births per 1,000 women aged 15–19) are from the United Nations Population Division, World Population Prospects. Unemployment rates are modeled ILO estimates from the International Labour Organization (ILOSTAT). All data correspond to 2005, except for lower secondary education completion rates, which are from 2004 for Chile and from 2003 for Argentina and Uruguay.

B Further Details about Data Construction

PANES/AFAM-PE records

Application and participation records contain application-form information from both successful and unsuccessful applications to PANES/AFAM-PE between April 2005 and December 2023. The information in these forms was collected through a detailed socio-demographic survey administered by program officials to applicant households. This is the data used to compute the poverty score that determines whether a household is eligible to participate in the program. The following table describes the raw variables included in these administrative records.

Description of Information Contained in Participation Records

Type	Variable
Application forms	Application number Resolution (accepted/rejected) Application date Self-reported income for each household member Poverty score (computed by the Social Security Agency)
Participation records	Monthly payment status and amount
Household Characteristics	Application Number Department Housing characteristics and quality Ownership status and property value Access to public services and utilities Appliances (microwave, fridge, car, computer, router, etc.)
Individual Characteristics	Application and masked national ID number Birth date Gender Education (current level and attendance) Activity and occupation status Income
Household Roster	Application and Masked ID numbers Relation with household head

Original *PANES* application records contain information of 187,727 application forms. Of these, 102,436 were accepted applications and 85,291 were rejected. In original *AFAM-PE* records, the total number of application forms is 796,131. This also includes households that were transferred automatically from *PANES* to *AFAM-PE*. Of the 796,131 *AFAM-PE* application forms, 585,052 (73.5%) were accepted and 211,079 (26.5%) were rejected. *AFAM-PE* individual level data contains information of about 1.6 million individuals.

Some individuals and application forms were removed from the data during the cleaning and preparation process. First, 6,150 application forms (0.62%) were removed due to missing information on household members, including missing individual ID numbers or cases where the forms could not be merged with the individual-level data. Then, 36 application forms were excluded because they had missing application dates, 472 because they lacked information on the poverty score, and 2,259 because they did not report the department in which they were submitted. Finally, there are several cases where multiple application forms for the same household are almost identical, including having the same application date, but differ in a few specific variables. In total, 7,840 forms are duplicated entries that differ only in the application form number, 1,030 differ only in the reason for rejection, and 4,921 report different poverty scores on the same application date. When the only difference is the application form number, I keep one form at random. When otherwise-identical forms differ only in the rejection reason, I collapse the reasons. For cases with multiple scores on the same day, I keep the form with the lowest score. The resulting number of application forms is 961,150, and the corresponding number of unique individuals is 1,711,140.

Birth Records

The birth certificate records come from the Ministry of Health and include the universe of births in the country (851,232) between 2003 and 2021. The raw dataset includes information on the birth date, the type of institution where the child was born (public, private, or other), the age of the mother at the time of birth, birth weight, and gestation weeks. The number of mothers in the dataset is 531,695. Of these, 57.3% gave birth to a single child during this period, 30.2% gave birth to two children, 9.0% to three children, and the remaining 3.5% to four or more children.

Starting in 2007, the Ministry of Health began transitioning birth certificates from paper to electronic format, which generated changes in the structure of the underlying databases. For this reason, the datasets can be organized into three periods: 2003–2007 (paper format), 2008–2010 (coexistence of paper and electronic records), and 2011 onward (fully electronic records, with the paper format aligned to the electronic structure). This means that while the key variables (such as the child's date of birth, the mother's ID number, birth weight, place of birth, and type of birth) are available throughout the entire period, other variables, mostly those related to the father and mother's pre-natal visits, are only available once the electronic records begin. In addition, there are changes over time in how marital and cohabitation status were reported.

Education records

The secondary education data come from administrative records from the National Administration of Public Education (ANEP) and cover lower and upper secondary education (grades 7–12) from 2004 to 2024. These records include information on enrollment, achievement, performance, and, when available (2015 and later), attendance. For grades 7–12, enrollment and achievement measures are consistently available throughout the entire period, while attendance is reported intermittently for lower secondary before 2016 and only reliably for upper secondary from 2022 onward. Coverage is nearly universal for public secondary schools, although the current version of the dataset does not yet include vocational education (UTU) or private schools. While vocational schools may be relevant for this lower-income population, private schools are not. Access to data on vocational schools has been authorized but, as of this writing, the data have not yet been provided. Data for private institutions are only available starting in 2023.

In total, the raw secondary education dataset includes 5,132,161 student-year observations for grades 7–12 between 2004 and 2024, corresponding to 1,138,501 unique students. Annual enrollment ranges from about 221,565 to 263,512 students per cohort. The distribution across grades in 2005, the year when *PANES/AFAM-PE* was rolled out was: approximately 19.5% of observations correspond to grade 7, 17.1% to grade 8, 16.6% percent to grade 9, 15.6% to grade 10, 18.1% to grade 11, and 13.1% percent to grade 12.

The raw data exhibit some heterogeneity in evaluation systems depending on the grade and study plan. In most cases, grades 7–10 follow a grade-level approval system, and the final qualification is available conditional on passing the grade. However, some study plans rely on course-by-course approval, and grades 11–12 are universally evaluated by subject without an aggregate grade.

Most students show up in records only once per academic year, although a small share have multiple inscriptions within the same year due to course-by-course evaluation structures or multiple inscriptions in different schools. Among students in grades 7–10, approximately 84.6% percent are evaluated under the grade-level approval system, while the remaining 15.4% percent follow course-by-course approval. For grades 11–12, all students are evaluated by subject.

To proxy for achievement outcomes that are comparable across evaluation systems, I use two measures that are well defined for all students in the sample: the maximum grade ever enrolled and an indicator for ever enrolling in 12th grade. These variables capture students' progression through the secondary education system and are consistently available across the full period. It is important to note that Uruguay does not count with a centralized record on secondary education completion, which would be the natural outcome to measure secondary education achievement.

Tertiary education data come from the *Universidad de la República*. In this case, the information provided includes enrollment records for the universe of students who ever enrolled in any major between 2005 and 2020. The only information currently available is whether the student was enrolled; no additional details on major, academic progress, performance, or completion status are available at this stage.

Labor market records

Labor market outcomes come from employer–employee matched administrative records from the Ministry of Labor and Social Security (MTSS). These data cover the universe of formal employment in Uruguay from January 2005 to July 2023 and include all registered workers (private and public), as well as registered self-employed, with at least one month of contributions to the social security system. The records provide monthly job-level information such as start and end dates, contract type, industry and firm characteristics, region, days worked, and wages, together with the date and reason for each separation (e.g., layoff, contract expiration, retirement). Although the data capture only formal employment, Uruguay's informal sector is comparatively small (about 17–20 percent of total employment).

In total, the dataset contains more than 269 million job–month observations corresponding to 2,688,860 unique workers during the 2005–2023 period. The month with the highest number of registered contracts is December 2022, with 1,365,509 jobs, while the lowest point occurs in January 2005, with 880,693. The data reflect a large expansion in formal employment between 2005 and 2015, during which registered jobs increased from about 880,000 to 1.3 million. This expansion was driven both by overall employment growth and by rising formalization (i.e., shifts from informal to formal work) during a period of sustained economic growth in Uruguay.

Household Identifiers

This section describes how I construct household identifiers used to define the reference application form. The goal is to assign individuals to the earliest application associated with their household history, going as far back as possible across the *PANES* and *AFAM-PE* records. This matters because household composition can change over time. For instance, a household may apply to *PANES* in 2005 with two parents and an older child, and later appear in *AFAM-PE* with the same parents, the older child, and a younger child born after the first application. Without linking these household records, the younger child would be assigned the later *AFAM-PE* application as the reference form, even though the household had already applied to *PANES*. The household crosswalk avoids this by identifying the earlier application associated with the household to which the child later belongs.

Unlike *AFAM-PE*, *PANES* records contain application form and personal identifiers but do not include a household identification number. A natural way to infer whether two applications correspond to the same household is to compare the personal identifiers listed in each form. If two forms include exactly the same individuals, they can be attributed to the same household. However, because household composition changes over time, for example because of births, separations, or household members moving, constructing a consistent household identifier requires a more careful approach.

I therefore construct harmonized household identifiers using strict rules based on shared household members. The procedure first creates household identifiers within the *PANES* records, then identifies candidate *PANES–AFAM-PE* household links, computes overlap-based matching rates, and classifies matches as perfect, good imperfect, or bad.

1. **Construct household identifiers within the *PANES* records.** I first create *PANES* household identifiers by linking the small share of application forms that have overlapping members. Forms are linked and assigned the same household identifier when they have identical or nested rosters, or when they satisfy overlap-based rules using shared individual identifiers, birth dates, names, and household-head information.[59] Forms with no repeated members are kept as unique households, while duplicated forms that do not satisfy these rules are kept as separate households.

    - This procedure assigns 187,727 *PANES* application forms to constructed household identifiers. The vast majority of forms are unique: 91.3% contain no listed member who appears in another form. Of the remaining forms, 6.3% are linked to another form through identical or very similar rosters, and 2.4% are classified as other cases.

2. **Identify candidate links between households in *PANES* and *AFAM-PE*.** I next create a pool of candidate links between *PANES* and *AFAM-PE* household identifiers that could correspond to the same underlying household. For this, I reconstruct household rosters in both datasets and identify all *PANES–AFAM-PE* household combinations that share at

[59] For non-identical forms, I classify two forms as belonging to the same household if one of the following conditions holds: they have the same household head and at least 50% of the origin-form members match; the household head differs or is missing, but at least 80% of members match in forms with more than one person; the comparison form is nested within the origin form and at least 60% of members match; or at least 60% of members match in two forms with at least three members.

least one individual ID. For *PANES*, I use the last observed roster for each constructed household. For *AFAM-PE*, I use the first observed roster for each household. This produces 252,168 candidate connections, corresponding to 137,391 unique *PANES* household IDs and 225,955 unique *AFAM-PE* household IDs. These represent 75.7% of the 181,559 constructed *PANES* household identifiers and 48.6% of the 464,559 *AFAM-PE* household identifiers.

3. **Identify overlapping household members.** For each candidate pair, I compare household rosters and identify overlapping members using individual ID numbers or, when ID numbers are missing, birth dates. This produces the person-level information used to compute household-pair matching rates.

4. **Compute household-pair match rates.** For each candidate pair, I define the matching rate as b/a, where b is the number of overlapping members identified by ID number or birth date, and a is the maximum household size across the two rosters. Pairs with a matching rate equal to one are classified as perfect matches; all other candidate pairs are classified as imperfect matches.

5. **Resolve cases with multiple perfect matches**. There are 176 cases where one *PANES* household has multiple perfect *AFAM-PE* matches, likely reflecting duplicate *AFAM-PE* household codes. In these cases, I assign the minimum *AFAM-PE* household ID. If a household has both perfect and imperfect matches, I retain only the perfect match.

6. **Drop clearly invalid imperfect matches.** I define these as household pairs with household size greater than or equal to three and a matching rate below 0.5, such as a four-person household where only one member overlaps.

7. **Classify good imperfect matches.** Among the remaining imperfect matches, I classify pairs as good matches if either:

   (a) the matching rate is at least 75%; or

   (b) the matching rate is lower than 75%, but the household head or partner matches and at least one child matches; or

   (c) both the household head and partner match, regardless of whether children match.

8. **Resolve multiple good matches.** For *PANES* households with multiple good imperfect matches, I select the *AFAM-PE* household with the highest matching rate.

In sum, this procedure mainly serves to recover household links across the two administrative systems and affects a limited set of ambiguous cases. Of the 181,559 constructed household identifiers in the *PANES* records, 137,391 share at least one individual with an *AFAM-PE* household, while 44,168 do not share any listed members. Among connected households, 73,878, or 53.8%, match exactly with at least one *AFAM-PE* household. The remaining cases require imperfect-match rules: 44,744 households are classified as good imperfect matches, while 18,767 are classified as bad matches and are not used to define household-level links.

Importantly, the constructed household links are used for a limited purpose: to identify the earliest reference application form associated with each individual's household history. This reference form determines the baseline score and household characteristics used in balance tests and as controls, and defines the household level used for clustering standard errors. The crosswalk is not used to construct the main outcomes. If anything, errors in this procedure would tend to introduce noise in the assignment of the reference application and attenuate discontinuities around the eligibility threshold, rather than mechanically generating the first-stage discontinuity documented in the main text.

	Perfect Matches				
Imperfect Matches	**0**	**1**	**2**	**3**	**Total**
0	0	39,823	118	1	39,942
1	32,015	19,870	45	0	51,930
2	16,742	8,496	8	0	25,246
3	8,068	3,421	2	0	11,491
4	3,822	1,309	1	0	5,132
5	1,652	496	0	0	2,148
6	717	194	1	0	912
7	308	73	0	0	381
8	114	15	0	0	129
9	42	5	0	0	47
10	17	2	0	0	19
11	8	0	0	0	8
12	3	0	0	0	3
13	2	0	0	0	2
14	1	0	0	0	1
Total	63,511	73,704	175	1	137,391

C Further Descriptive Statistics on Baseline Characteristics and Outcomes

In this appendix, I describe in more detail the characteristics of the samples used in the analysis, focusing on individual and household characteristics reported in the application forms, as well as summary statistics for the main outcomes of interest.

C.1 Baseline Characteristics

Figure C.1 shows the age distribution of individuals in the main sample at two key points in time. Panel (a) shows individuals' age at the time of their household's first application to *PANES/AFAM-PE*, while Panel (b) focuses on the distribution of age on 31 December 2021. Individuals in the *main sample* have ages at first application that range between 2.8 and 17.8 years old, with an average of 9.6, and a median of 9.8. Overall, the figure shows that there is significant variation in the timing of exposure. Panel (b) shows that individuals are on average 26 years old in December 2021, when fertility outcomes are last observed. This also mean that they are on average 27.4 years old in July 2023, when labor market outcomes are last observed, with ages ranging between 20.6 to 36.6.

Table C.1 reports descriptive statistics for the main individual and household characteristics observed in the *PANES/AFAM-PE* application forms. Columns (1) and (2) present these statistics for the universe of applicants, without any restriction. Columns (3) and (4) focus on the *full sample*, which, as defined in Section 2.2, refers to individuals who were younger than 18 years old at the time of their household's first application and at least 19 years old in December 2021. Columns (5) and (6) focus on the *main sample*, which applies further restrictions to the *full sample* related to the years of application (2005–2007) and outlier application behavior. Odd-numbered columns report descriptives for individuals across the whole support of the poverty score, whereas even-numbered columns focus on individuals who are close to the eligibility threshold. For simplicity, the bandwidth used for this table is 0.02, which corresponds to the maximum optimal bandwidth reported in Table 2.

In terms of individual characteristics, individuals in the *main sample* who are close to the eligibility threshold, described in column (6), were on average 10.2 years old at the time of their household's first application to the program and are about 26.3 years old in December 2021. On average, they appear in 3.0 application forms throughout the period. This includes applications in which they were dependent children, as well as potential applications later on as household heads or household head partners in newly formed households. The sample is split roughly evenly between men and women (51.4 percent vs. 48.6 percent). Individuals in this group look similar to those in the full support of the main sample, described in (5), and to individuals in the full sample, described (3) and (4). The main difference is the number of applications, which is higher in the *main sample* because these individuals applied earlier and therefore had a longer window to appear in subsequent applications. Compared to the universe of applicants, which also includes adults and is described in (1) and (2), their age is mechanically lower since the focus here is on children living in applicant households.

When looking at the characteristics of the first application form, individuals in the *main* *sample* close to the threshold have an average standardized score of roughly 0.0 and an average first-application acceptance rate of 57.4%. Only about 15.5% of these individuals correspond to application forms submitted in the capital city, a share that is considerably lower than the roughly 37% of the population living in the capital. Compared to the universe of applicants or to the *full sample* close to the threshold, individuals in the *main sample* are more likely to be accepted in their first application (57.3% vs. 49.3% and 54.2%, respectively) and less likely to come from the capital city (15.5% vs 23.0% and 19.8%, respectively). By design, individuals close to the threshold have scores clustered around zero and lower acceptance rates, as the overall distribution of the poverty score is heavily skewed toward households above the eligibility cutoff.

Finally, in terms of household-level characteristics, individuals in the *main sample* close to the threshold come from households with an average of 4.3 members, an average member age of 23.1, and 52.9% of them being single-parent households. Looking at household head characteristics, 62.6% of heads report being employed (either formally or informally), and they have on average 6.5 years of education, which is slightly above completion of primary school. The average household head income is USD 130.8, which is substantially below the poverty line (approximately USD 350). Overall, there are no substantial differences in household characteristics when comparing the *main sample* to the *full sample* or to the universe of applicants, except, by construction, in the share of applications submitted before September 2005. Differences between households close to the threshold and those in the full support remain similar across samples.

Overall, the main takeaway is that the *main sample* looks very similar to the *full sample* of children along most observable characteristics, including age at application, gender, household size, household head education, and household head labor income. The main differences arise mechanically from the restrictions imposed on the main sample: individuals in this group applied earlier (2005-2007) and therefore appear in more application forms. In addition, households in the *main sample* are slightly less likely to be located in the capital city. Apart from these differences, the *main sample* resembles the broader population of children growing up in applicant households.

C.2 Participation in *PANES/AFAM-PE* and Other Outcomes of Interest

In addition, Table C.2 reports descriptive statistics for the main outcomes of interest used throughout the analysis for the *full* and *main samples*. Statistics are shown both for the full support of the poverty score and for individuals whose first application was within 2 p.p. of the threshold, and separately for men and women.

Panel (a) focuses on *PANES/AFAM-PE* participation outcomes: total amount collected from the program before age 18, months in the program, and an indicator for ever participating. On average, individuals in the *main sample* close to the threshold collected around USD 9,000 from the program, spent between 72 and 74 months in the program, and 87% participated for at least one month, with no differences between men and women.

Panel (b) reports descriptive statistics for labor market outcomes. In the *main sample* close to the threshold, 56.1% of women were employed for at least one month during the last year (i.e., August 2022-July 2023). Over the same period, women worked an average of 4.8 months and earned around USD 3,760, or USD 7,000 among those with positive earnings. The average tenure in the current job, conditional on having one, is 39 months. In cumulative terms, 70% of women have ever been employed for at least four consecutive months in the formal sector. On average, women accumulated 37 months of formal work, earned around USD 24,600, and had 2.7 different employers. As expected, men in the same sample exhibit higher labor market attachment: 61.2% were employed during the last year, working an average of 5.2 months and earning around USD 4,580, or USD 7,940 among those with positive earnings. For men with a current job, average tenure is 42 months. Cumulatively, 77% of men have ever been employed. On average, men accumulated 46 months of formal work, earned around USD 33,600, and had 3.7 different employers.

Panel (c) reports descriptive statistics for outcomes related to individuals' transition to adulthood. In particular, the table shows that 23.6% of women in the *main sample* close to the threshold had a career-oriented transition. I define someone as taking a career-oriented path if two conditions hold: (i) they did not have a child before age 19, and (ii) they either had their first relatively stable job (defined as an employment spell of at least four consecutive months) before age 19 or enrolled in university. When looking at the different dimensions that characterize this transition, Table C.2 shows that the average age at first birth for these women is 20.0 years, and 21.1% had a birth before age 19. Regarding employment-related transition outcomes, the average age at first employment is 20.7 years, and 17.4% had a spell of at least four consecutive months in a formal job before age 19. Finally, only 11.0% of women in this sample ever enrolled in university. This suggests that tertiary education enrollment is not the most relevant decision margin for *PANES/AFAM-PE* applicants. Men in the same sample exhibit a higher prevalence of career-oriented transitions, at 33.9%. Their average age at first birth is 22.8 years, and only 2.2% appear in birth records as becoming fathers before age 19. The average age at first employment is 19.6 years, and 31.4% had a job before age 19. University enrollment is even lower than for women, at 4.5%, which is consistent with general patterns of men being less attached to the education system than women.

Panel (d) reports descriptive statistics for fertility outcomes. In the *main sample* close to the threshold, 55.3% of women had at least one birth by the end of the observation period. On average, women in this sample had 0.92 children. Men in the same sample display substantially lower fertility rates in the administrative records: 23.4% are recorded as having had a child, with an average of 0.31 children. These differences likely reflect both biological timing and differences in the probability that fathers are linked to birth records.

Panel (e) reports descriptive statistics for secondary education outcomes. In the *main sample* close to the threshold, 80.0% of women ever enrolled in secondary school. On average, they spent 4.0 years enrolled and reached a maximum grade of 10.17. In addition, 28.7% of women ever enrolled in Grade 12. Men in the same sample show lower educational attachment: 64.7% ever enrolled in secondary school, they spent 2.7 years enrolled on average, and reached a maximum grade of 9.32. Only 12.1% ever enrolled in Grade 12. These differences indicate substantially lower progression within secondary education among men relative to women.

Panel (f) reports descriptive statistics for adult welfare program applications and participation. These outcomes correspond to *PANES/AFAM-PE* applications made after childhood, when individuals could apply as household heads or partners of the household head in newly formed households. In the *main sample* close to the threshold, 35.0% of women applied to the program as adults, and their average age at first adult application was 21.8 years. Moreover, 32.4% of women were accepted into the program, and the average woman received 21 payments. Adult participation is substantially lower among men: 16.8% applied as adults, with an average age at first adult application of 23.2 years, while 15.7% were accepted and received on average 9.1 payments. These differences reflect both lower application and acceptance rates among men.

Overall, the patterns described in each panel broadly hold when considering the full support of the poverty score within the *main sample* and when looking at the *full sample*. The main differences are in magnitudes, following two consistent patterns. First, when expanding to the full support of the poverty score, outcomes generally look worse: individuals collected more benefits as children, show weaker labor market outcomes, less career-oriented transitions, earlier fertility, display lower educational attainment, and higher adult welfare participation. This reflects the fact that the full support includes individuals farther away from the threshold, and the distribution of the poverty score is skewed toward the eligible region, that is, poorer households. Second, when considering the *full sample*, outcomes look somewhat better on average. This likely reflects that households who applied later, as the program expanded, were in relatively better socioeconomic conditions than those covered at the beginning.

Figure C.1: Distribution of Age at 31 December, 2021 and at First Application

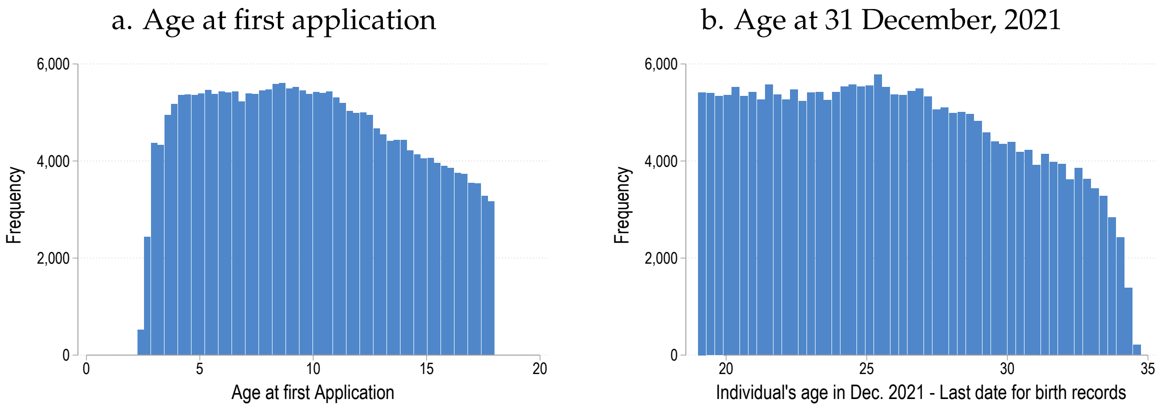

Notes: This figure illustrates the age distribution of individuals in the *main sample* at two key moments in their lives. Panel (a) depicts the age distribution at the time of the first application, while Panel (b) shows the age distribution in December 2021, the last moment for which birth certificates are available. As described in Section 3.1, to compute the age at first application, I first identify the earliest application submitted by each household. I then assign to each individual the earliest of these household-level first applications for the households in which they appear as members on the corresponding form.

Table C.1: Descriptive Statistics: Individual Characteristics

	Universe		Full Sample		Main Sample	
	Full Support (1)	Bdw: 0.02 (2)	Full Support (3)	Bdw: 0.02 (4)	Full Support (5)	Bdw: 0.02 (6)
a. Individual Characteristics						
Female (%)	53.17	53.84	49.70	49.62	48.50	48.63
	(49.90)	(49.85)	(50.00)	(50.00)	(49.98)	(49.98)
Avg. Number of Applications	2.41	2.47	2.85	2.82	3.24	3.04
	(1.48)	(1.46)	(1.56)	(1.50)	(1.40)	(1.37)
Age of 1st Application	20.30	22.95	10.69	10.68	9.84	10.16
	(20.68)	(21.22)	(4.13)	(4.19)	(4.22)	(4.26)
Age at 31 Dec. 2021	32.72	36.35	25.49	25.91	26.02	26.31
	(21.05)	(21.60)	(4.03)	(4.16)	(4.21)	(4.25)
b. Reference Form						
Std. Score	0.15	0.00	0.17	0.00	0.16	-0.00
	(0.23)	(0.01)	(0.24)	(0.01)	(0.21)	(0.01)
App. Accepted (%)	65.92	49.28	71.25	54.22	76.02	57.35
	(47.40)	(50.00)	(45.26)	(49.82)	(42.70)	(49.46)
Capital City (%)	34.50	22.95	31.81	19.81	27.06	15.50
	(47.54)	(42.05)	(46.57)	(39.86)	(44.43)	(36.19)
c. Household characteristics (ref. form)						
Number of Members	4.10	3.67	4.74	4.17	4.94	4.26
	(1.76)	(1.64)	(1.91)	(1.79)	(1.99)	(1.88)
Single Parent (%)	43.84	50.18	45.40	52.07	45.28	52.91
	(49.62)	(50.00)	(49.79)	(49.96)	(49.78)	(49.92)
Mother's Age When 1st Born	23.38	24.48	23.01	24.11	22.56	23.94
	(7.52)	(7.81)	(7.01)	(7.31)	(6.62)	(7.14)
Avg. age	23.62	26.07	21.80	23.54	20.98	23.14
	(11.39)	(11.70)	(7.40)	(7.93)	(7.21)	(8.00)
Baseline App. Before Sep.2005	18.58	23.08	28.31	30.36	41.07	38.75
	(38.89)	(42.13)	(45.05)	(45.98)	(49.20)	(48.72)
Household Head: Employed (%)	56.70	59.37	61.65	61.67	64.58	62.63
	(49.55)	(49.11)	(48.62)	(48.62)	(47.83)	(48.38)
Household Head: Ed. years	6.91	6.94	6.58	6.79	6.30	6.56
	(2.81)	(2.88)	(2.59)	(2.61)	(2.35)	(2.44)
Household Head: income	163.19	155.65	155.28	155.36	121.30	130.84
	(250.29)	(215.64)	(212.59)	(187.65)	(135.00)	(131.90)
Observations	1,711,132	128,021	420,383	32,771	248,693	25,030

Notes: This table reports a series of descriptive statistics for the *universe* of PANES/AFAM-PE applicants, individuals in the *full sample*, and individuals in the *main sample*. As defined in Section 2.2, the universe of *PANES/AFAM-PE* applicants corresponds to all individuals who appear in *PANES/AFAM-PE* application forms from the start of the program in April 2005 until December 2023. The *full sample* is defined as individuals who were younger than 18 years old when their parents first applied to PANES/AFAM-PE and at least 19 years old by December 2021 (N=420,383). The *main sample* is the sample for which I implement the RDD analysis and corresponds to the full sample after a series of restrictions explained in detail in the main text (N=248,693). Odd-numbered columns include summary statistics for individuals across the full support of the poverty score obtained in the first application (Z^{1st}), while even-numbered columns are based on individuals who are close to the eligibility threshold. For simplicity, I use a fixed threshold of 0.02, which corresponds to the maximum optimal bandwidth used in estimates for the main outcomes reported in Table 2. Panel (a) reports information on a series of characteristics at the individual level. Panel (b) focuses on the characteristics of the reference form, i.e., the application form corresponding to the first application. Finally, panel (c) reports information about the characteristics of the household defined in the first application form.

Table C.2: Descriptive Statistics: Outcomes of Interest

	Full Sample				Main Sample			
	Full Support		Opt. Bdw.		Full Support		Opt. Bdw.	
	Women	Men	Women	Men	Women	Men	Women	Men
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. *PANES/AFAM-PE* Participation Outcomes, before 18yo								
Total Transfer Collected (1000s, USD)	9.40	9.40	7.87	7.89	11.67	11.43	9.08	9.01
	(7.55)	(7.43)	(6.92)	(6.88)	(7.59)	(7.50)	(7.02)	(7.02)
Months in the *PANES/AFAM-PE*	73.88	73.76	65.64	65.49	88.25	86.27	74.01	72.88
	(55.17)	(54.24)	(53.57)	(52.96)	(56.01)	(55.26)	(54.36)	(53.83)
Ever Treated before 18yo (%)	88.21	87.86	84.15	84.17	92.01	91.85	87.08	87.09
	(32.25)	(32.66)	(36.53)	(36.50)	(27.11)	(27.35)	(33.55)	(33.53)
b. Labor Market Outcomes								
Last Observed: Employed in the last year	55.59	60.86	58.63	62.75	52.49	58.25	56.09	61.23
	(49.69)	(48.81)	(49.25)	(48.35)	(49.94)	(49.31)	(49.63)	(48.73)
Last Observed: Months Worked in the last year	4.69	5.13	5.13	5.41	4.31	4.77	4.84	5.17
	(5.11)	(5.11)	(5.21)	(5.15)	(5.01)	(5.04)	(5.17)	(5.11)
Last Observed: Total Earnings (1000s, USD)	3.53	4.47	4.03	4.80	3.16	4.13	3.76	4.58
	(5.05)	(5.72)	(5.43)	(5.90)	(4.83)	(5.59)	(5.32)	(5.83)
Last Observed: Total Earnings (1000s, USD), (>0)	6.61	7.80	7.21	8.18	6.25	7.48	7.03	7.94
	(6.21)	(7.31)	(6.56)	(7.76)	(6.05)	(7.02)	(6.49)	(7.25)
Last Observed: Tenure in current job	34.44	38.68	37.90	40.93	35.62	40.04	38.91	42.36
	(33.18)	(37.07)	(35.73)	(38.61)	(35.07)	(39.19)	(37.09)	(40.38)
Cumulative: Ever Employed (%)	67.80	74.77	71.37	77.24	66.00	74.02	69.97	76.95
	(46.72)	(43.43)	(45.21)	(41.93)	(47.37)	(43.85)	(45.84)	(42.12)
Cumulative: Months Worked	32.84	41.24	37.41	45.52	32.11	41.88	37.03	46.41
	(39.33)	(43.49)	(42.12)	(45.94)	(40.19)	(45.07)	(42.96)	(47.28)
Cumulative: Total Earnings (1000s, USD)	21.51	30.04	25.27	33.43	20.43	29.74	24.62	33.58
	(31.82)	(37.36)	(34.95)	(39.37)	(31.68)	(37.75)	(35.08)	(39.83)
Cumulative: Number of Different Employers	2.54	3.41	2.70	3.58	2.48	3.56	2.66	3.71
	(2.72)	(3.21)	(2.75)	(3.23)	(2.73)	(3.33)	(2.76)	(3.32)
c. Transition Outcomes								
Market Oriented (%)	25.17	33.90	27.74	35.90	20.17	31.13	23.55	33.93
	(43.40)	(47.34)	(44.77)	(47.97)	(40.13)	(46.30)	(42.43)	(47.35)
Age at 1st Birth	19.65	22.34	20.02	22.65	19.56	22.50	19.98	22.75
	(3.17)	(3.32)	(3.39)	(3.41)	(3.22)	(3.39)	(3.42)	(3.45)
Birth Before 19 (%)	21.16	2.12	18.75	2.04	25.00	2.32	21.13	2.19
	(40.85)	(14.40)	(39.03)	(14.12)	(43.30)	(15.06)	(40.82)	(14.65)
Age at 1st Employment	20.42	19.56	20.57	19.57	20.64	19.66	20.70	19.64
	(2.77)	(2.32)	(2.85)	(2.38)	(2.92)	(2.44)	(2.93)	(2.48)
Job Before 19 (%)	18.73	30.47	18.71	31.89	16.97	29.47	17.35	31.41
	(39.02)	(46.03)	(39.00)	(46.61)	(37.54)	(45.59)	(37.87)	(46.42)
Enrolled in University (%)	11.60	5.61	14.43	6.40	7.69	3.31	10.96	4.52
	(32.02)	(23.02)	(35.14)	(24.47)	(26.65)	(17.90)	(31.24)	(20.77)
d. Fertility Outcomes								
Had a Birth (%)	52.16	19.66	50.11	20.99	58.91	22.37	55.25	23.39
	(49.95)	(39.74)	(50.00)	(40.73)	(49.20)	(41.67)	(49.73)	(42.33)
Number of Births	0.87	0.25	0.82	0.27	1.02	0.29	0.92	0.31
	(1.05)	(0.57)	(1.02)	(0.59)	(1.11)	(0.61)	(1.05)	(0.62)
e. Secondary Education Outcomes								
Ever Enrolled in Secondary School (%)	78.43	65.19	80.79	67.21	77.10	61.56	80.00	64.74
	(41.13)	(47.64)	(39.39)	(46.95)	(42.02)	(48.64)	(40.00)	(47.78)
Years Enrolled in Secondary School	3.69	2.63	3.97	2.79	3.62	2.40	3.97	2.65
	(3.00)	(2.74)	(3.03)	(2.80)	(3.04)	(2.68)	(3.09)	(2.78)
Ever Enrolled in Grade 12 (%)	27.71	13.18	33.12	15.51	21.79	9.08	28.74	12.06
	(44.76)	(33.82)	(47.07)	(36.20)	(41.28)	(28.73)	(45.26)	(32.56)
Max. Grade Enrolled in Secondary	10.09	9.39	10.36	9.56	9.79	9.06	10.17	9.32
	(1.77)	(1.78)	(1.69)	(1.77)	(1.79)	(1.73)	(1.73)	(1.75)
f. Welfare Outcomes								
Applied as Adult (%)	34.47	15.59	31.07	14.94	40.49	18.43	35.01	16.82
	(47.53)	(36.27)	(46.28)	(35.65)	(49.09)	(38.78)	(47.70)	(37.40)
Age at 1st Application as Adult	21.47	22.85	21.82	23.18	21.46	22.92	21.82	23.20
	(3.59)	(3.61)	(3.78)	(3.70)	(3.66)	(3.68)	(3.84)	(3.75)
Accepted as Adult (%)	32.23	14.61	28.53	13.86	38.04	17.30	32.37	15.65
	(46.74)	(35.32)	(45.16)	(34.55)	(48.55)	(37.82)	(46.79)	(36.33)
Number of Payments Received	20.31	8.27	18.04	7.86	25.40	10.28	21.08	9.11
	(38.64)	(25.32)	(37.19)	(24.89)	(42.81)	(28.41)	(39.82)	(26.80)
Std. Poverty Score	0.25	0.32	0.23	0.29	0.26	0.33	0.23	0.30
	(0.19)	(0.18)	(0.19)	(0.17)	(0.19)	(0.18)	(0.19)	(0.17)
Observations	208,932	211,451	16,261	16,510	120,610	128,083	12,173	12,857

Notes: This table reports a series of descriptive statistics for individuals in the *full* and *main sample*. As defined in Section 2.2, the *full sample* consists of individuals who were younger than 18 years old when their parents first applied to PANES/AFAM-PE and at least 19 years old by December 2021 (N=420,383). The *main sample* is the sample for which I implement the RDD analysis and corresponds to the full sample after a series of restrictions explained in detail in the main text (N=248,709). Columns (1) and (2) present separate statistics for men and women in the *full sample* across the full support of the poverty score obtained in the first application (Z^{1st}), while columns (3) and (4) do the same for individuals who are close to the eligibility threshold. For simplicity, I use a fixed threshold of 0.02, which corresponds to the maximum optimal bandwidth used in estimates for the main outcomes reported in Table 2. Columns (5) through (8) are analogous for individuals in the *main sample*. The variables included in this table correspond to outcomes that are typically of interest. Panel (a) focuses on variables related to participation in *PANES/AFAM-PE*, Panel (b) focuses on labor market outcomes, Panel (c) on outcomes related to transitions to adulthood, Panel (d) to fertility outcomes, Panel (e) on outcomes from secondary education records, and Panel (f) on downstream outcomes related to participation in *PANES/AFAM-PE* as adults.

D Further Details on the Research Design

In this appendix I provide additional details about the validity of the RDD.

Application of eligibility rule from the perspective of the program's administrators. I begin by documenting how program administrators applied the eligibility rule. Since the goal is to report whether application forms were effectively accepted or rejected based on the poverty score obtained, the unit of analysis is the application form, regardless of the household or the household members. Figure D.1 illustrates this by describing the relationship between the standardized poverty score and the likelihood that an application form was accepted. Panel (a) shows this relationship over the full support of the score, while Panel (b) focuses on forms within the optimal bandwidth. In both panels, the figure reports average acceptance rates across bins of the poverty score and shows an extremely sharp discontinuity at the eligibility threshold.

Despite this sharp discontinuity, the score rule was not perfectly applied on either side of the threshold. Three patterns are worth noting. First, to the right of the threshold, the acceptance rate was lower than 100%. This can be explained by other rejection reasons unrelated with the poverty score. For instance, the program also required households to pass an income test at the moment of application. As reported in Amarante et al. (2013), in practice, only 10% of applicant households were effectively rejected because of this reason. While I am not able to test this directly in the data, since *PANES* records I have access to do not contain information on rejection reason, this income-related rejection rate is consistent with the 90% acceptance rate observed to the right of the threshold in Figure D.1. Second, to the left of the threshold, the acceptance rate being higher than zero is most likely explained by moving eligibility thresholds over time and automatic enrollment during the transition from *PANES* to *AFAM-PE*. In both cases, however, the size of non-compliance is small, and the probability of acceptance still changes sharply at the cutoff.

Finally, it is worth noting that the acceptance rate for some ineligible forms extremely close to the threshold (within 0.0015, or 0.15 p.p.) was abnormally high. This is most likely due to precision issues in the raw data, which in some cases were provided as string variables with only four decimal places. For this reason, throughout the analysis the baseline specification excludes application forms within a radius of 0.0015 of the eligibility threshold. When excluding these observations, Panel (b) documents that the size of the discontinuity is 76 p.p. (p-value $<$ 0.001), reflecting a jump from around a 17% acceptance rate just below the threshold to about 93% just above it. In any case, for first-stage and main outcomes I also report how sensitive the estimates are to the choice of the donut radius. Overall, these tests show that all main results hold regardless of the specific radius choice.

First Stage and Distribution of Z^{1st} across the whole support. Figure D.2, panel (a), expands the analysis of the relationship between the amount of *PANES/AFAM-PE* benefits collected before age 18 and the running variable shown in panel (a) of Figure 1 to the full support of Z^{1st}. Two main patterns emerge from this figure. First, there is a sharp discontinuity at the eligibility threshold, consistent with the one reported in the main text. Second, when considering the full support of Z^{1st}, there is a clear positive relationship between the poverty score at first application and the total amount collected from the cash transfer. This indicates that households in worse conditions were more exposed to the benefits of the program. This pattern may reflect earlier entry into the program, longer participation (i.e., greater difficulty in moving out of poverty and becoming ineligible), higher amounts due to larger number of children, or a combination of these factors. Panel (b) complements this description by plotting the distribution of Z^{1st} across the full support. The figure shows that most applications lie in the eligible region of the threshold, but there is still a substantial share of individuals whose households were ineligible at the time of first application. This is critical for the use of a RDD approach.

Distribution of amount collected and months in *PANES/AFAM-PE*. Figure D.3 describes the distribution of two variables related to participation in *PANES/AFAM-PE*. Panel (a) shows the distribution of the total amount collected from *PANES/AFAM-PE* before age 18, measured in USD 1,000. Overall, the figure shows that most individuals in the *main sample* participated in *PANES/AFAM-PE* for at least some period. This is reflected in the fact that only about 20,000 individuals collected between USD 0 and USD 1,000. In addition, the figure illustrates substantial variation in the total amount collected: most observations lie between USD 1,000 and USD 40,000, a wide range. The figure also depicts the 95th percentile, which is the threshold used to winsorize the cumulative amount collected before age 18, that is, the endogenous variable used in most specifications. Panel (b) replicates this exercise using the number of months in the program instead of the total amount collected. In this case, the variable is capped at 192 months, which corresponds to 16 years of participation. This includes households who started receiving benefits during the *PANES* phase (2005–2007), as well as those who entered the program when *PANES* was rebranded as *AFAM-PE* and have remained in the program since then.

Alternative interpretations for the discontinuity at the threshold. To provide alternative interpretations of the discontinuity observed at the eligibility threshold, Figure D.4 replicates the first stage illustration in panel (a) of Figure 1 in the main text using two alternative measures of *PANES/AFAM-PE* participation. First, panel (a) depicts the discontinuity in the probability of *ever* participating in *PANES/AFAM-PE*, defined as a binary variable (0,100) for collecting benefits for at least one month. Panel (a) shows a discontinuity of 18.4 p.p. (24.4%) at the eligibility threshold. This difference is statistically significant at conventional levels (p-value ≤ 0.001). Second, panel (b) shows the same figure using months in the program as the outcome of interest. A similar discontinuity is observed: eligible individuals were exposed to the program before age 18 for, on average, 20 additional months (32.7%) compared to individuals just below the threshold.

Robustness and sensitivity test for first-stage estimates. As described in Section 4.1, a natural concern in RDD settings is whether the observed discontinuity at the eligibility threshold reflects a genuine jump in the outcome or instead arises from specific modeling choices or sampling variation. To address this concern, I implement four robustness and sensitivity exercises: (i) a placebo cutoff test using a dense grid of alternative thresholds, (ii) sensitivity to bandwidth choice, (iii) robustness to alternative RDD technical specifications, and (iv) variation in the donut radius around the cutoff. These tests are reported in Figure D.5, panels (a) through (d), respectively.

Implementing these tests to cumulative *PANES/AFAM-PE* transfers collected before age 18 rule out that the discontinuities observed in the baseline specification are driven by specific and arbitrary choices in the RDD. The placebo cutoff exercise produces a distribution of estimates centered around zero, and the estimate at the true eligibility threshold is, by far, the largest in absolute value. Furthermore, the estimated discontinuity remains positive and statistically significant across a wide range of manually selected bandwidths and results are also stable under alternative RDD specifications, including different kernel functions, polynomial degrees, and bandwidth selectors. Finally, varying the donut radius does not materially affect the direction or statistical significance of the first stage. However, as expected, the size of the discontinuity attenuates when including observations in a 0.0015 range of the eligibility threshold. Precisely because they can be thought of as a source of measurement error that weakens the relevance of the first stage, these observations are excluded in the baseline specification. Overall, these exercises support the interpretation of the observed jump in cumulative transfers as a genuine discontinuity at the eligibility threshold.

Trends in Extensive Margin of Participation in *PANES/AFAM-PE* Section 3.1 discusses time trends in household income variables, including labor income and transfers. In particular, panel (c) of Figure 1 shows a large jump in household transfers at the time of the first application for eligible children, with effects that fade out after 2–3 years. Figure D.6 confirms that this fade-out is driven mostly by ineligible households entering the program during the transition from *PANES* to *AFAM-PE*, when the eligibility threshold became more lenient (i.e., the control group catching up in participation). In particular, the share of households just below the threshold receiving *PANES/AFAM-PE* increases sharply from about 20% in the year of first application to roughly 80% in the third year after application, and declines thereafter. Since around 80% of individuals in the *main sample* lived in households that applied for the first time in 2005, the third year after first application corresponds mainly to 2008. This is precisely the year in which *PANES* transitioned into *AFAM-PE*, when the eligibility threshold became more lenient and there was a substantial increase in the number of new participants.

Discontinuity in Predicted Eligibility Finally, Figure D.7, depicts a summary test for systematic differences in baseline characteristics just at the threshold. More specifically, I compute the predicted eligibility score based on a probit model that uses first-time eligibility ($\mathbb{1}(Z_i^{1st} > 0)$) as the dependent variable and all other observed baseline characteristics as the independent variables. The fact that there is no discontinuity in predicted eligibility at the threshold ($\tau_{SRD} = -0.002$, $p - value = 0.740$) can be interpreted as a signal of no systematic discontinuities in the baseline variables.

Figure D.1: Relation Between Application Form Eligibility and Resoultion

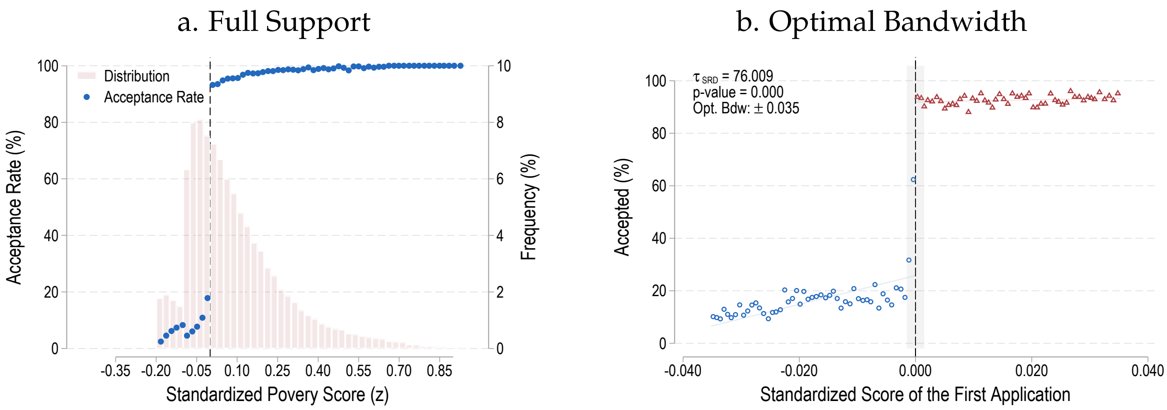

Notes: This figure reports the share of applications that were accepted into the program as a function of the standardized poverty score (Z). Each data point used to construct this figure corresponds to an application form. The sample of interest corresponds to all application forms associated with individuals in the *main sample* as defined in Section 3. Each application form is included only once. Panel (a) shows this relation for the full support of Z. Negative values of Z indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest. The number of bins was selected manually such that the number of bins for negative values of Z relative to the number of bins for the positive values of Z represents the distribution of z. Panel (b) focuses on application forms that are located within an optimal bandwidth. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Outcomes are grouped into 50 quantile-spaced bins on each side of the threshold. In addition, the figure reports the point estimate of the local difference in the outcome of interest at the threshold (τ_{SRD}), based on Equation (3), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. Inference is based on robust standard errors clustered at the household level.

Figure D.2: First Stage and Distribution of the Poverty Score for the Full Support

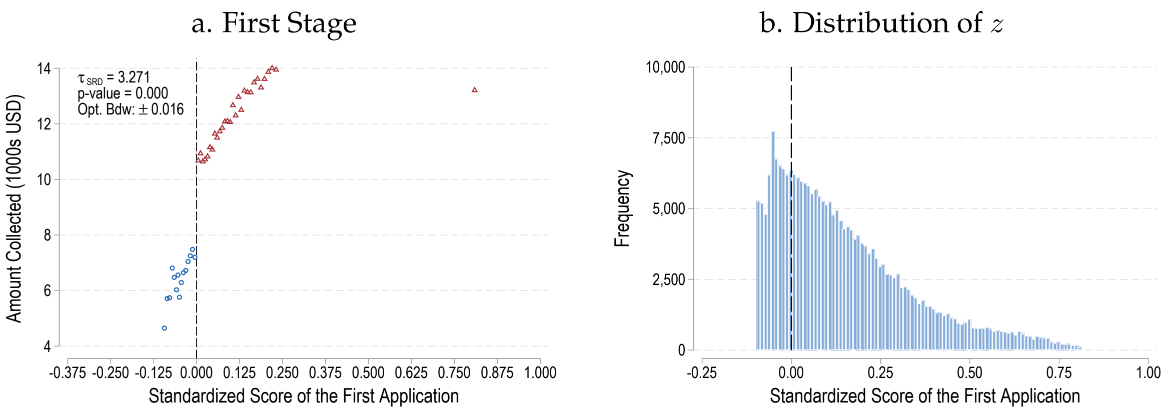

Notes: This figure expands the analysis of the validity of the RDD depicted in Figure 1 in the main text to the full support of the standardized poverty score at first application (Z_i^{1st}). Panel (a) shows the average total benefits received from *PANES/AFAM-PE* before age eighteen, expressed in 2008 USD (1000s), as a function of the standardized poverty score at first application (Z_i^{1st}), over the full support of Z_i^{1st}, rather than only the optimal bandwidth as in Panel (a) of Figure 1. Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest; after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. The number of bins is selected manually such that the number of bins for negative values of Z_i^{1st} relative to the number of bins for positive values of Z_i^{1st} reflects the distribution of Z_i^{1st}. While the figure is depicted over the full support, the RDD estimates reported in the upper-left corner are based on the optimal bandwidth. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. The figure reports the point estimate of the local difference in the outcome of interest at the threshold (τ_{SRD}), based on Equation (3), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. The sample corresponds to the *main sample* as defined in Section 3, excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (b) reports the distribution of Z_i^{1st} for the *main sample*, in steps of 1%. For visual convenience, the figure excludes 5% (about 10,000 cases) of individuals whose non-standardized poverty scores lie within 0.001 of 0 or 1.

Figure D.3: Distribution of Months and Amount Collected from *PANES/AFAM-PE*

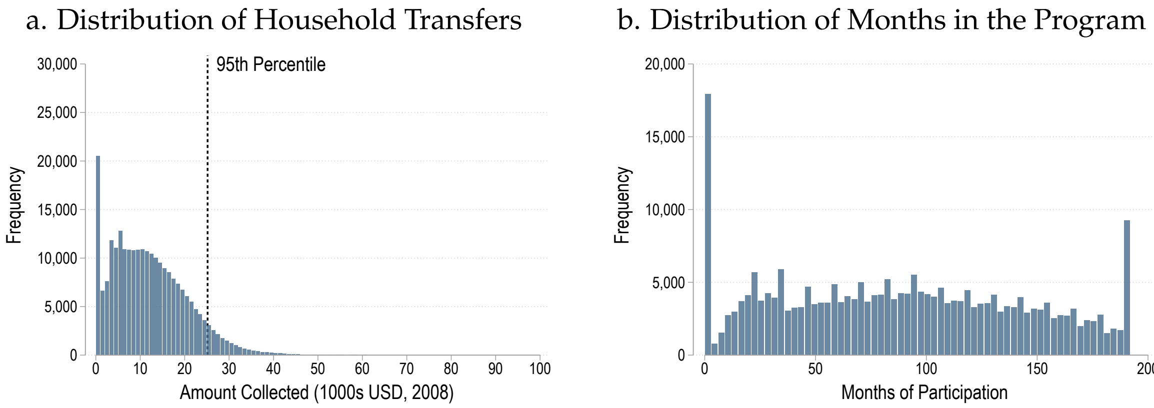

Notes: This figure reports the distribution of program exposure for individuals in the *main sample*. Panel (a) shows the distribution of the total amount of benefits received from *PANES/AFAM-PE* before age eighteen, expressed in 2008 USD (1000s). The vertical line in Panel (a) indicates the 95th percentile of the distribution, which is used for winsorization across specifications. Panel (b) shows the distribution of cumulative months of participation in the program before age eighteen. The vertical line in Panel (b) indicates 16 years of participation. This corresponds to households who started receiving benefits during the *PANES* phase (2005–2007), as well as those who entered the program when *PANES* was rebranded as *AFAM-PE* and have remained in the program since then. The sample used corresponds to the *main sample* as defined in Section 3, excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach).

Figure D.4: First Stage Estimates for Alternative Definitions of the Endogenous Variable

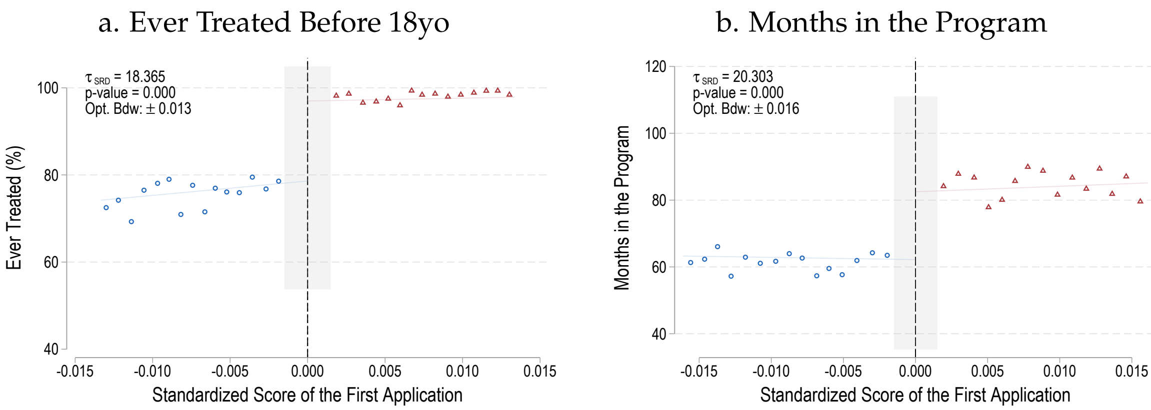

Notes: This figure reports a series of alternative first-stage RDD estimates for individuals in the *main sample*, excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Each panel illustrates the outcome of interest as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest, after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Average total benefits are grouped into 20 quantile-spaced bins on each side of the threshold (i.e., about four times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). In addition, the figure reports the point estimate of the local difference in average total benefits at the threshold (τ_{SRD}), based on Equation (2), together with the optimal bandwidth and the continuity test p-value, shown in the upper-left corner. Panel (a) reports the probability of ever receiving *PANES/AFAM-PE* benefits before age eighteen. This is defined as a binary variable (0, 100) that equals 100 if the individual ever belonged to a household that received at least one month of *PANES/AFAM-PE* transfers prior to turning eighteen, and 0 otherwise. Panel (b) reports cumulative months of participation in the program before age eighteen. Inference is based on robust standard errors clustered at the household level.

Figure D.5: Robustness Tests: First Stage

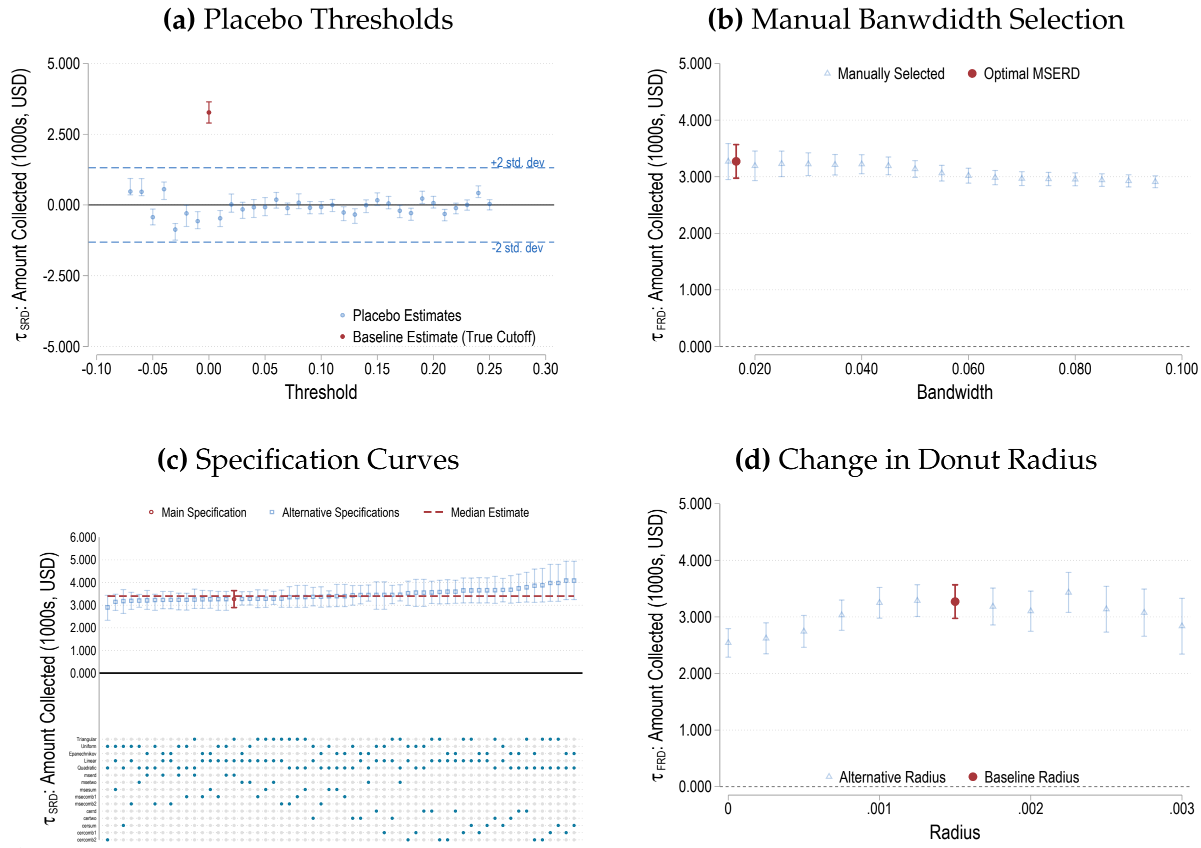

Notes: This figure reports a series of robustness checks for the first-stage RDD estimates using the *main sample*. In all cases, the outcome of interest is average total benefits received from *PANES/AFAM-PE* before age eighteen, expressed in 2008 USD (1000s). All estimates take as the starting point the first-stage specification reported in Figure 1, Panel (a), which is described in detail in the figure notes, and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots different τ_{SRD} obtained under these alternative specifications together with the corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the first-stage regression at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{SRD} obtained by estimating the first-stage regression using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{SRD} obtained by estimating the first-stage regression under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{SRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{SRD} obtained by estimating the first-stage regression under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level.

Figure D.6: Trends in Extensive Margin of Participation in *PANES/AFAM-PE*

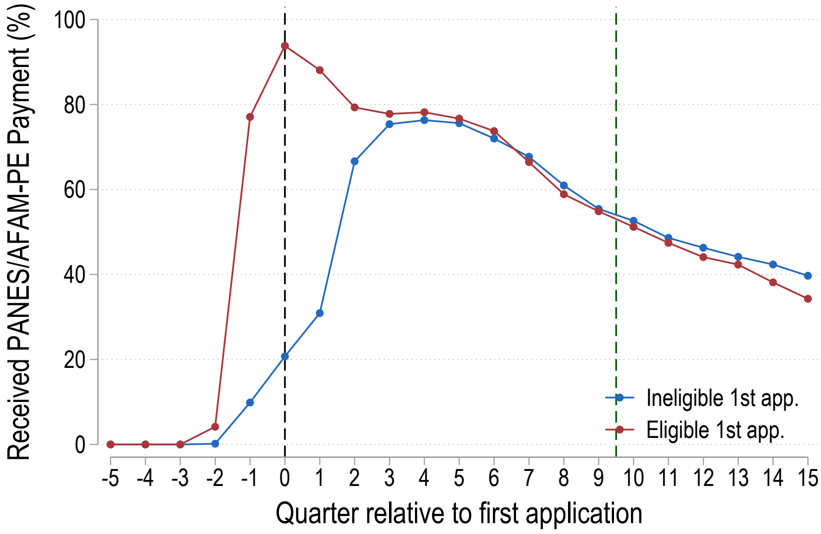

Notes: This figure illustrates the time dynamics of extensive-margin participation in *PANES/AFAM-PE* close to the eligibility threshold. The analysis is conducted for individuals in the *main sample*, excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). As in Figure 1, panel (c), the horizontal axis reports years relative to the first application date, with year 0 corresponding to the year of first application. For each year, a first-stage RDD regression is estimated where the dependent variable is a binary indicator equal to 100 if the individual belongs to a household that receives at least one *PANES/AFAM-PE* transfer in that year, and 0 otherwise. The specification follows the extensive-margin first-stage estimates reported in Figure D.4, Panel (a). Each bin reported in the figure corresponds to the lower or upper intercept of the local linear regression underlying the RDD estimation. For illustrative purposes, these estimates are obtained without additional covariates, so that each intercept can be interpreted as the local average participation rate on each side of the threshold for a given year.

Figure D.7: Discontinuity in Predicted Eligibility

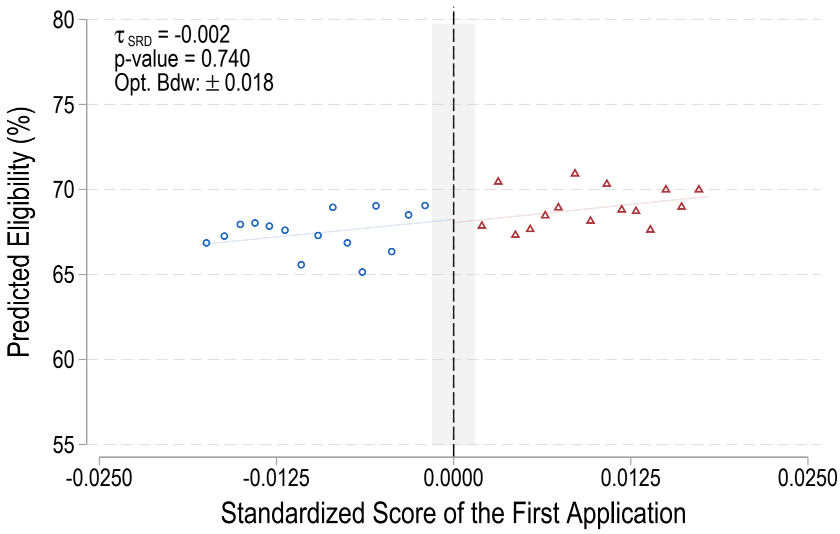

Notes: This Figure reports a regression discontinuity–based summary test for systematic differences in baseline characteristics at the eligibility threshold. The figure is based on the *main sample*, excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). The outcome is the predicted probability of eligibility, obtained from a probit model that uses first-time eligibility ($\mathbb{1}(Z_i^{1st} > 0)$) as the dependent variable and all other observed baseline characteristics $\mathbf{X}$ as regressors. The figure illustrates this predicted eligibility score as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the predicted eligibility score, after adjusting for the set of fixed effects $\mathbf{\Lambda}$ described in Section 3.1. In this specification, I omit $\mathbf{X}$ as these are the variables used to compute the predicted eligibility score. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. The outcome is grouped into 20 quantile-spaced bins on each side of the threshold (i.e., about four times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). The figure also reports the point estimate of the local difference in predicted eligibility at the threshold (τ_{SRD}), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. Inference is based on robust standard errors clustered at the household level. Analogous continuity tests for each baseline characteristic considered individually are reported in Table 1.

E Further Results on Labor Market Outcomes

In this Appendix I show a series of additional results that complement the main analysis on labor market outcomes discussed in Section 4.1

Sensitivity of main estimates to the inclusion of additional covariates and fixed effects. The baseline specification used to compute the effects of additional *PANES/AFAM-PE* support on labor market outcomes is described in detail in Equation (1). In addition to the standard components of an RDD, this specification includes two sets of additional variables, denoted by $\mathbf{\Lambda}$ and $\mathbf{X}$. The vector $\mathbf{\Lambda}$ corresponds to year of birth, year of application, and region fixed effects. These fixed effects are included to rely on more comparable units, for example within region or within cohort. The vector $\mathbf{X}$ represents baseline characteristics reported in the application forms. The inclusion of $\mathbf{X}$ is not motivated by evidence of baseline discontinuities that would threaten the identification assumption, as these are ruled out by Table 1 in the main text. Instead, these variables are included to increase statistical precision. The set of variables in $\mathbf{X}$ coincides with those reported in Table 1.

Tables E.1 and E.2 assess the sensitivity of the baseline results to the inclusion of these additional variables. First, Table E.1 replicates the specification in Table 2, including fixed effects $\mathbf{\Lambda}$ but excluding baseline controls $\mathbf{X}$. Second, Table E.2 further removes the fixed effects $\mathbf{\Lambda}$. In both cases, to facilitate comparison with the baseline specification, both the bandwidth and the bias correction are set equal to those used in Table 2.

Across both alternative specifications, the direction of the estimates remains unchanged. For women, the results continue to show positive effects on cumulative labor market outcomes and weaker responses on last-year outcomes, independently of the inclusion of additional controls. For men, the results also remain unchanged, as there is no evidence of economically meaningful effects on their labor market outcomes across specifications, except for some positive effects on last-year months worked and earnings. As discussed in Section 4.1 and examined in more detail below, these latter estimates are highly sensitive to the bandwidth choice and other technical aspects of the RDD specification.

More specifically, when removing only the baseline covariates $\mathbf{X}$ (i.e., retaining fixed effects $\mathbf{\Lambda}$), magnitudes and statistical significance remain essentially unchanged relative to the baseline specification. If anything, point estimates are slightly smaller. For example, the 2SLS estimate for women's cumulative months worked is 2.2 compared to 2.5 in the baseline specification. Similarly, the estimate for cumulative earnings is now 1.6 compared to 1.7 in the baseline specification. Larger differences arise when fixed effects $\mathbf{\Lambda}$ are also excluded. In this case, the point estimates for cumulative outcomes decrease to 1.9 (months worked) and 1.4 (earnings). More importantly, standard errors increase by 22% and 16%, respectively. As a consequence, statistical significance falls to the 5% and 10% levels, compared to 1% in the baseline specification.

Taken together, these results show that the inclusion of baseline covariates $\mathbf{X}$ and fixed effects $\mathbf{\Lambda}$ in the baseline specification is not driving the main findings. The additional controls mostly improve precision, particularly through the inclusion of fixed effects, but do not alter the direction or gender pattern of the estimated effects substantially. For this reason, the baseline specification includes both sets of variables, as they help increasing efficiency.

Sensitivity of main estimates to winsorizing. Table E.3 reports estimates analogous to those in Table 2, but using non-winsorized variables. This applies both to the endogenous variable, i.e., cumulative transfers, and to the earnings-related outcome variables: last-year earnings and cumulative earnings. The results show almost no differences between the non-winsorized and winsorized estimates, with very similar results in terms of direction, magnitude, and statistical significance across all outcomes.

Baseline robustness and sensitivity tests for main outcomes. As discussed in Section 4.1 and Appendix D, a standard concern in RDD settings is whether the estimated discontinuity reflects a true jump at the eligibility threshold rather than modeling choices or sampling variation. In the main text, Figure 3 reports a set of robustness tests for women's cumulative earnings: (i) placebo cutoffs over a dense grid of alternative thresholds, (ii) alternative bandwidth choices, (iii) alternative RDD technical specifications, and (iv) variation in the donut radius around the cutoff (panels (a) through (d), respectively). Here, I replicate the same four exercises for each column in Table 2. Figures E.1 to E.4 report the results for women (last-year months worked, last-year earnings, cumulative months worked, and cumulative earnings), and Figures E.5 to E.8 present the corresponding results for men. For conciseness, I do not describe each figure separately and instead focus on the overall patterns.

For women, Figures E.1 to E.4 yield two main conclusions. First, and most importantly, estimates for cumulative outcomes, reported in Figures E.3 and E.4, are highly robust across all sensitivity exercises. The effects remain positive and statistically significant across a wide range of bandwidths (from 0.01 to 0.095), alternative kernel and bandwidth selection procedures, different local polynomial specifications, and variations in the donut radius. As expected, magnitudes decline as the bandwidth increases, reflecting the standard bias–variance trade-off when moving away from very local comparisons in an RDD setting. However, even when the optimal bandwidth is multiplied by more than six, estimates remain positive and statistically significant at the 10% level. In the placebo cutoff tests, estimates at the true eligibility threshold are the largest among all alternative thresholds and the only ones exceeding two standard deviations from the mean placebo estimate, which is centered around zero. Finally, conditional on the optimal bandwidth, neither the specific RDD specification nor the choice of donut radius pushes the estimates toward the extremes of the distribution; the baseline estimate lies close to the median across specifications.

When examining the robustness tests for last-year outcomes, reported in Figures E.1 and E.2, the conclusions remain consistent with those obtained under the baseline specification. Across sensitivity exercises, the direction of the effects is stable under alternative bandwidth choices, RDD specifications, and donut sizes. However, because last-year effects are substantially smaller than cumulative effects, widening the bandwidth leads to a loss of statistical significance, particularly for months worked. As discussed in Section 4.1, the smaller and less precisely estimated effects for last-year outcomes likely reflect both the additional variability associated with a single-year measure and the dynamics of treatment effects over the life cycle. Since last-year outcomes are measured at older ages, they may reflect later-life realizations of treatment effects rather than the full accumulation of labor market responses.

When examining robustness tests for men's outcomes, two main patterns emerge. First, for last-year outcomes, reported in Figures E.5 and E.6, the sensitivity to bandwidth choice is similar to that observed for women. As the bandwidth increases, magnitudes attenuate and statistical significance weakens. However, in contrast to women's estimates, the results for men are substantially more sensitive to specific technical choices in the RDD specification. In particular, positive and statistically significant effects are primarily obtained under linear local regressions and when bandwidths are selected using MSE-optimal procedures. When instead using quadratic local regressions or CER-optimal bandwidth selectors, the corresponding estimates are generally smaller and statistically indistinguishable from zero.

Second, and more importantly, cumulative effects oscillate between positive, null, and even negative values depending on the bandwidth selected, as shown in Figures E.7 and E.8. At the optimal bandwidth, and across alternative kernel functions, polynomial degrees, and donut sizes, the estimated effects remain close to zero and statistically insignificant under the baseline specification. However, when varying the bandwidth manually, the estimates become negative under several specifications and, in some cases, statistically significant when using bandwidths larger than approximately 0.04–0.05. Given that cumulative outcomes aggregate labor market behavior over the full early-adulthood period and constitute the less volatile measure of treatment effects, the absence of stable positive cumulative effects, jointly with their sensitivity in sign under alternative specifications and the strong sensitivity of last-year estimates to modeling decisions, supports an interpretation of null effects for men.

Estimates under alternative definitions of participation in *PANES/AFAM-PE*. In a fuzzy RDD framework, the definition of the endogenous variable mostly affects the scaling of τ_{FRD} and, therefore, its interpretation. Alternative definitions of the endogenous variable rely on the same eligibility-induced discontinuity for identification (see Section 3.1 for a more detailed discussion). However, when comparing estimates across specifications, it is also important to keep in mind that changes in the definition of the endogenous variable can also affect the optimal bandwidth selected by the MSERD criterion, so that the estimation sample may slightly differ across specifications.

Because the main goal of this paper is to study how additional income support during childhood affects future life trajectories, the most natural endogenous variable is cumulative transfers received before age 18. This measure captures total monetary exposure induced by the program and allows the reduced-form discontinuity at the threshold to be scaled by the first stage, yielding 2SLS estimates expressed in monetary units (e.g., per additional USD 1000 of income support), under linearity assumptions. However, alternative re-scaling approaches may also be informative. Tables E.4 and E.5 replicate estimates in Table 2 for two alternative definitions of participation.

First, instead of total transfers, Table E.4 reports estimates using months of participation in the program as the endogenous variable. In this case, the 2SLS coefficients should be interpreted as the effect of one additional month of exposure to *PANES/AFAM-PE* before turning 18 years old. For instance, consider the estimated effect on women's cumulative earnings reported in column (4). One additional month of exposure to *PANES/AFAM-PE* increases cumulative earnings by

USD 265 (1.08%, p-value = 0.002). The interpretation of the first stage also changes relative to the baseline specification and now reflects the reduced-form discontinuity in months of exposure at the eligibility threshold. In column (4), the first-stage coefficient of 19.1 implies that women just above the eligibility threshold, based on their first application score, participate in *PANES/AFAM-PE* for approximately 19 additional months on average. This magnitude is consistent with the patterns documented in Figure D.6, which shows that after *PANES* was expanded into *AFAM-PE*, roughly two years after the initial implementation of *PANES*, most applicants who were initially ineligible but close to the threshold entered the program. Importantly, differences in reduced-form estimates across specifications are mechanically associated with changes in the optimal bandwidth selected by `rdrobust`. Conditional on fixing the bandwidth (and bias correction), reduced-form estimates should not change when alternative definitions of the endogenous variable are used.

Similarly, Table E.5 replicates the baseline specification using an indicator for ever participating in *PANES/AFAM-PE* before age 18 as the endogenous variable. In this case, τ_{FRD} corresponds to the effect of moving from no participation to ever participating, under a linearity assumption (i.e., each additional percentage point increase in the probability of participating in *PANES/AFAM-PE* has the same marginal effect on the outcome of interest). Following with the previous example, consider the estimated effect on women's cumulative earnings reported in column (4). The 2SLS estimate of 20.8 implies that moving from no participation to ever participating in *PANES/AFAM-PE* increases cumulative earnings by USD 20,805 (83%, p-value = 0.001). The interpretation of the first stage also changes relative to the baseline specification. In this case, the first-stage coefficient of 0.19 reflects the reduced-form discontinuity in the probability of ever participating at the eligibility threshold. This implies that individuals just above the threshold are 18.9 percentage points more likely to have ever participated in *PANES/AFAM-PE* before age 18 (24.8%).

At this point, it is useful to clarify the interpretation of these magnitudes. When the endogenous variable is defined as an indicator for ever participating, the 2SLS estimator recovers the Local Average Treatment Effect (LATE) at the threshold. The first-stage discontinuity of 0.19 reflects the share of compliers, that is, individuals whose participation status is affected by eligibility. The estimated coefficient therefore captures the effect of a discrete change from non-participation to participation for this subgroup.

Estimates on additional labor market outcomes. In Table 2, I described the effects of additional income support on last-year months employed and last-year earnings. Tables E.6 and E.7 extend this analysis by reporting effects on additional labor market outcomes for women and men, respectively. These capture additional dimensions of labor market attachment and earnings. Columns (1) through (4) correspond to binary indicators (0–100) of labor market attachment: ever employed (at least one spell of four consecutive months), employed at least one month in the last year, employed all 12 months in the last year, and holding a job by July 2023. Columns (5) reports estimates on log-wages received in July 2023, while column (6) report effects on log earnings between August 2022 and July 2023. In both cases, by construction, the effects are conditional on being employed in July 2023, and employed at least once between August 2022-July 2023, respectively.

For women, Table E.6, column (1), replicates the baseline 2SLS, reduced-form, and first-stage estimates for the ever-employed variable. The 2SLS coefficient is -0.372 (-0.52%, p-value = 0.667), implying that an additional USD 1,000 of childhood income support does not affect the probability of ever being employed in the formal sector. The corresponding reduced-form discontinuity is -1.2 p.p. (-1.69%, p-value = 0.666).

Columns (2) through (4) report alternative measures of last-year employment attachment. The null effect persists when considering the probability of being employed for at least one month during the last year. However, columns (3) and (4) show positive and statistically significant effects on other measures of last-year employment attachment. In column (4), which examines the probability of being employed in July 2023 (the last month observed in the data), the 2SLS estimate indicates that an additional USD 1,000 of childhood income support increases this probability by 1.4 p.p. (3.4%, p-value = 0.057), with a corresponding reduced-form effect of 4.4 p.p. (10.7%, p-value = 0.054). When considering a stronger measure of labor market attachment, i.e., the probability of being employed for the full 12 months of the last year, the effects are larger and statistically significant at the 1% level. The 2SLS estimate implies that an additional USD 1,000 of childhood income support increases the probability of holding a job during all 12 months of the previous year by 2.1 p.p. (8.1%, p-value = 0.007), with a reduced-form effect of 6.9 p.p. (26.6%, p-value = 0.005). These patterns suggest that, although last-year effects are smaller than cumulative effects, they are concentrated on more stable forms of last-year employment. I discuss this more in depth in Section 4.2 of the main text, where I analyze the effects of additional childhood income support on job mobility and later-life labor market stability in more detail.

Columns (5) and (6) report log-earnings outcomes conditional on employment. In both cases, the estimates are positive but not statistically significant at conventional levels. Conditional on employment in July 2023, the 2SLS estimate implies an increase of approximately 1% in wages (p-value = 0.182), with a corresponding reduced-form effect of 3.1% (p-value = 0.173). Similarly, column (6) reports a 2SLS estimate of 2.8% for cumulative earnings in the last year (p-value = 0.181), with a reduced-form discontinuity of approximately 9% (p-value = 0.173). Although these effects are imprecisely estimated, their positive sign is consistent with the broader pattern documented in Section 4.1 and with the robustness evidence presented in Figures E.1 and E.2, which shows that last-year effects are smaller and estimated with less precision than cumulative outcomes.

Table E.7 presents the corresponding estimates for men. For the probability of ever being employed the effects are null and both statistically and economically insignificant. For last-year employment, whether is any number of months or the full year, the estimated effects are positive, close to 1 p.p. (between 1.5% and 4.0%), and marginally insignificant (p-values = 0.138 and 0.126). The probability of being employed in July 2023 shows a 1.9p.p. effect, which is statistically significant effect at the 5% level (p-value = 0.012). Turning to earnings, the patterns are mixed. The effect on log wages in July 2023 is -1.5% (p-value = 0.240), while the effect on last-year log earnings is 4.7% and statistically significant at the 10% level (p-value = 0.083). Overall, these weak and inconsistent effects on last-year outcomes are in line with the main results reported in Table 2. As discussed extensively in Section 4.1 and in this appendix, robustness and sensitivity tests show that these estimates are highly dependent on bandwidth choice and technical RDD specifications. Combined with cumulative effects that oscillate between positive, null, and negative values depending on the specification, this mixed pattern reinforces an interpretation closer to null effects for men.

Effects on Firms' Characteristics. To examine whether additional childhood income support affected the characteristics of the firms in which individuals are employed, Tables E.8 and E.9 replicate the baseline strategy using a series of firm characteristics measured in July 2023 as outcomes for women and men, respectively. These estimates are conditional on being employed at that time. Column (1) reports effects on the log number of employees, column (2) on log total wage bill, column (3) on log median wage, and column (4) on log firm age. In general, the overall evidence does not indicate systematic differences in the characteristics of the firms in which individuals are employed, For number of employees, total wage bill, and firm age, the estimated effects for women are null and economically small. The only partial exception is median wage, where the estimated effect is positive and approximately 1.6%. However, this estimate is statistically insignificant (p-value = 0.205). For men, the results broadly replicate the patterns observed for women. The only exception is the coefficient on firm age, which indicates a positive effect of approximately 5%, statistically significant at the 10% level (p-value = 0.075). Despite this, the overall evidence suggests that additional childhood income support did not meaningfully alter the characteristics of the firms in which individuals are employed. However, estimates should be interpreted cautiously as they are conditional on being employed in July 2023, and, as shown in Table E.6, additional income support increases the probability of being employed at that time, particularly for women.

While estimates on firm quality cannot be easily interpreted because of selection, an alternative way to examine where labor market gains are concentrated is to analyze the unconditional effects on employment by sector. Figure reports the τ_{FRD} coefficients from separate regressions for dummy variables corresponding to 18 activity sector for women in the *main sample*. The binary variables are defined as 1 if an individual is employed in a given sector and 0 otherwise, which includes employment in other sectors as well as not being employed. These estimates reflect both changes in sector allocation among marginal workers induced to enter formal employment and shifts in the distribution of sectors among those who would have been employed in any case. As such, the interpretation is descriptive and focuses on where increases in formal employment are concentrated. For women, the unconditional sectoral results reveal clear employment gains that are concentrated in retail trade, with moderate increases in manufacturing and administrative services and no strong offsetting declines in other sectors. These patterns are consistent with the overall positive effects on formal employment being driven primarily by increased employment in these private-sector industries.

When looking at men, the picture is, as in previous outcomes, more mixed. Figure shows that while there are increases in employment in manufacturing and transportation, these are accompanied by a relatively large decline in public administration and smaller declines in agriculture and construction. Unlike for women, and consistent with the effects discussed in the previous section, these results are more aligned with a story of null effects on employment and earnings for men. One possible interpretation is that men's responses to additional childhood income support are more closely related to reallocation across sectors and less to increased or more sustained participation in the labor market. In any case, given the absence of strong effects on cumulative labor market outcomes, as already documented and discussed in detail, even if men reallocated across sectors, this has not yet translated into substantial gains in earnings or accumulated experience. Furthermore, as shown in Appendix F, effects are also null for men when examining mobility and employment stability outcomes.

Table E.1: Effects on Labor Market Outcomes - Including FE, but no Controls, Baseline Bandwidth and Bias

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed (1)	Earnings (2)	Months Employed (3)	Earnings (4)	Months Employed (5)	Earnings (6)	Months Employed (7)	Earnings (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.099	0.163*	2.239***	1.585***	0.160*	0.159*	0.589	0.336
	(0.084)	(0.088)	(0.651)	(0.561)	(0.070)	(0.081)	(0.527)	(0.469)
Robust p-value	0.274	0.086	0.001	0.008	0.057	0.090	0.300	0.521
Mean Baseline Outcome	4.89	3.80	37.16	24.59	5.18	4.61	47.33	34.28
Effect Size (%)	2.03%	4.29%	6.02%	6.45%	3.08%	3.45%	1.24%	0.98%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.310	0.511*	6.985***	4.928***	0.519**	0.520*	1.915	1.093
	(0.258)	(0.266)	(1.895)	(1.659)	(0.222)	(0.260)	(1.704)	(1.520)
Robust p-value	0.264	0.074	0.000	0.005	0.043	0.074	0.282	0.506
Mean Baseline Outcome	4.89	3.80	37.16	24.59	5.18	4.61	47.33	34.28
Effect Size (%)	6.34%	13.43%	18.80%	20.04%	10.03%	11.27%	4.05%	3.19%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.125***	3.128***	3.120***	3.109***	3.253***	3.265***	3.253***	3.254***
	(0.273)	(0.277)	(0.283)	(0.289)	(0.240)	(0.249)	(0.240)	(0.241)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.49	7.52	7.55	7.55	7.40	7.39	7.41	7.40
Effect Size (%)	41.71%	41.61%	41.32%	41.15%	43.94%	44.17%	43.91%	43.95%
Selection of Bandwidth:								
Opt. Bandwidth	[0.017]	[0.017]	[0.016]	[0.016]	[0.019]	[0.018]	[0.019]	[0.019]
Effective Obs.	9,203	9,066	8,755	8,549	11,006	10,391	11,017	10,930

Notes: This table replicates Table 2 in the main text using a specification *with* fixed effects $\mathbf{\Lambda}$ but *without* additional control variables $\mathbf{X}$. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), except for not including $\mathbf{X}$ as controls, interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3), excluding $\mathbf{X}$. Panel (c) reports first-stage estimates, δ, based on Equation (2), also excluding $\mathbf{X}$. To facilitate comparison with the baseline specification, both the bandwidth and the bias correction are set equal to those used in Table 2. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.2: Effects on Labor Market Outcomes - No FE, no Controls, Baseline Bandwidth and Bias

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.086	0.147	1.889**	1.367*	0.122	0.111	-0.005	-0.093
	(0.082)	(0.089)	(0.780)	(0.641)	(0.065)	(0.075)	(0.567)	(0.475)
Robust p-value	0.342	0.129	0.024	0.051	0.145	0.235	0.965	0.815
Mean Baseline Outcome	4.89	3.80	37.16	24.59	5.18	4.61	47.33	34.28
Effect Size (%)	1.76%	3.85%	5.08%	5.56%	2.36%	2.41%	-0.01%	-0.27%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.284	0.483	6.167***	4.421**	0.440	0.403	-0.018	-0.334
	(0.264)	(0.276)	(2.229)	(1.864)	(0.226)	(0.266)	(2.041)	(1.719)
Robust p-value	0.320	0.101	0.009	0.026	0.111	0.200	0.965	0.810
Mean Baseline Outcome	4.89	3.80	37.16	24.59	5.18	4.61	47.33	34.28
Effect Size (%)	5.80%	12.69%	16.60%	17.98%	8.50%	8.74%	-0.04%	-0.98%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.301***	3.293***	3.264***	3.234***	3.599***	3.621***	3.598***	3.599***
	(0.353)	(0.357)	(0.366)	(0.372)	(0.312)	(0.323)	(0.312)	(0.313)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.49	7.52	7.55	7.55	7.40	7.39	7.41	7.40
Effect Size (%)	44.06%	43.81%	43.23%	42.82%	48.60%	48.99%	48.57%	48.61%
Selection of Bandwidth:								
Opt. Bandwidth	[0.017]	[0.017]	[0.016]	[0.016]	[0.019]	[0.018]	[0.019]	[0.019]
Effective Obs.	9,203	9,066	8,755	8,549	11,006	10,391	11,017	10,930

Notes: This table replicates Table 2 in the main text, but uses a specification without fixed effects $\mathbf{\Lambda}$ and without additional control variables $\mathbf{X}$. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), except for not including $\mathbf{\Lambda}$ or $\mathbf{X}$ as controls, interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3), excluding $\mathbf{\Lambda}$ and $\mathbf{X}$. Panel (c) reports first-stage estimates, δ, based on Equation (2), also excluding $\mathbf{\Lambda}$ and $\mathbf{X}$. To facilitate comparison with the baseline specification, both the bandwidth and the bias correction are set equal to those used in Table 2. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.3: Effects on Labor Market Outcomes - Non-Winsorized Variables

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.124	0.190**	2.301***	1.305**	0.152**	0.155	0.582	0.288
	(0.076)	(0.092)	(0.567)	(0.525)	(0.068)	(0.095)	(0.519)	(0.642)
Robust p-value	0.101	0.039	0.000	0.011	0.027	0.116	0.190	0.458
Mean Baseline Outcome	4.92	4.03	37.30	27.09	5.25	4.95	47.86	39.62
Effect Size (%)	2.51%	4.72%	6.17%	4.82%	2.90%	3.13%	1.22%	0.73%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.415*	0.646**	7.726***	4.385***	0.507**	0.516	1.935	0.959
	(0.254)	(0.305)	(1.791)	(1.736)	(0.222)	(0.314)	(1.716)	(2.133)
Robust p-value	0.097	0.036	0.000	0.010	0.023	0.108	0.183	0.454
Mean Baseline Outcome	4.92	4.03	37.30	27.09	5.25	4.95	47.86	39.62
Effect Size (%)	8.44%	16.01%	20.71%	16.19%	9.66%	10.43%	4.04%	2.42%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.362***	3.391***	3.358***	3.361***	3.325***	3.333***	3.326***	3.326***
	(0.291)	(0.302)	(0.289)	(0.254)	(0.240)	(0.251)	(0.239)	(0.246)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.61	7.63	7.62	7.69	7.52	7.49	7.52	7.52
Effect Size (%)	44.20%	44.44%	44.05%	43.73%	44.24%	44.50%	44.20%	44.26%
Selection of Bandwidth:								
Opt. Bandwidth	[0.018]	[0.018]	[0.019]	[0.022]	[0.020]	[0.019]	[0.021]	[0.020]
Effective Obs.	8,830	8,355	8,907	10,886	10,452	9,651	10,519	10,025

Notes: This table replicates Table 2 in the main text, but removes winsorization both from the endogenous variable as well as from the earnings-related variables. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3), corresponding to the sharp discontinuities shown in Figures 2 through 4. Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.4: Effects on Labor Market Outcomes - End. Var.: Months in *PANES/AFAM-PE*

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Months Collected	0.012	0.019*	0.411***	0.265***	0.025**	0.020	0.092	0.030
	(0.011)	(0.011)	(0.104)	(0.084)	(0.012)	(0.014)	(0.091)	(0.081)
Robust p-value	0.251	0.090	0.000	0.002	0.044	0.163	0.280	0.665
Mean Baseline Outcome	4.92	3.86	37.21	24.56	5.24	4.66	47.86	34.86
Effect Size (%)	0.23%	0.48%	1.10%	1.08%	0.47%	0.43%	0.19%	0.09%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.223	0.360*	7.949***	5.066***	0.521**	0.422	1.946	0.633
	(0.205)	(0.218)	(1.841)	(1.526)	(0.237)	(0.286)	(1.927)	(1.713)
Robust p-value	0.246	0.086	0.000	0.001	0.038	0.157	0.273	0.663
Mean Baseline Outcome	4.92	3.86	37.21	24.56	5.24	4.66	47.86	34.86
Effect Size (%)	4.52%	9.33%	21.36%	20.63%	9.95%	9.06%	4.07%	1.81%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	19.320***	19.272***	19.363***	19.119***	21.013***	21.245***	21.210***	21.230***
	(1.453)	(1.513)	(1.853)	(1.763)	(1.642)	(1.736)	(1.720)	(1.728)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	63.21	63.42	63.23	63.64	61.38	61.68	61.68	61.66
Effect Size (%)	30.56%	30.39%	30.62%	30.04%	34.24%	34.45%	34.39%	34.43%
Selection of Bandwidth:								
Opt. Bandwidth	[0.025]	[0.024]	[0.018]	[0.019]	[0.018]	[0.017]	[0.017]	[0.017]
Effective Obs.	12,448	11,641	8,522	9,204	9,345	8,560	8,660	8,600

Notes: This table replicates Table 2 in the main text, but uses cumulative *months of participation* in *PANES/AFAM-PE* as the endogenous variable in the fuzzy RDD specification. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), but this time using the endogenous variable is cumulative months of participation in *PANES/AFAM-PE*. In this case, estimates should be interpreted as the effect of an additional month of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2), where the first stage relates eligibility to cumulative months of program participation. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.5: Effects on Labor Market Outcomes - End. Var.: Ever Treated

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Ever in *PANES/AFAM-PE*	0.539	1.113	29.652***	20.805***	0.345	-0.049	-4.681	-5.236
	(0.739)	(0.784)	(7.337)	(6.095)	(0.657)	(0.683)	(4.719)	(4.192)
Robust p-value	0.436	0.127	0.000	0.001	0.465	0.938	0.471	0.323
Mean Baseline Outcome	4.92	3.88	37.54	25.09	5.46	5.05	49.89	36.80
Effect Size (%)	10.95%	28.69%	78.98%	82.92%	6.33%	-0.97%	-9.38%	-14.23%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.105	0.217	5.586***	3.940***	0.064	-0.009	-0.869	-0.972
	(0.144)	(0.152)	(1.345)	(1.134)	(0.122)	(0.127)	(0.876)	(0.778)
Robust p-value	0.439	0.128	0.000	0.001	0.471	0.938	0.474	0.325
Mean Baseline Outcome	4.92	3.88	37.54	25.09	5.46	5.05	49.89	36.80
Effect Size (%)	2.13%	5.59%	14.88%	15.70%	1.18%	-0.18%	-1.74%	-2.64%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	0.194***	0.195***	0.188***	0.189***	0.186***	0.186***	0.186***	0.186***
	(0.008)	(0.009)	(0.011)	(0.011)	(0.007)	(0.007)	(0.007)	(0.007)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	0.76	0.76	0.77	0.76	0.73	0.73	0.73	0.73
Effect Size (%)	25.70%	25.63%	24.57%	24.77%	25.41%	25.44%	25.41%	25.43%
Selection of Bandwidth:								
Opt. Bandwidth	[0.045]	[0.043]	[0.029]	[0.030]	[0.055]	[0.065]	[0.064]	[0.065]
Effective Obs.	22,614	21,301	14,303	15,113	30,610	36,241	35,676	36,166

Notes: This table replicates Table 2 in the main text, but uses an indicator for ever receiving *PANES/AFAM-PE* during childhood as the endogenous variable in the fuzzy RDD specification. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), where the endogenous variable is an indicator for ever receiving *PANES/AFAM-PE*. In this case, estimates are interpreted as the effect of receiving childhood income support (relative to never receiving it) on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2), where the first stage relates eligibility to the probability of ever receiving *PANES/AFAM-PE*. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.6: Effects on Additional Labor Market Outcomes - Women

	Labor Market Attachment				Earnings Cond. on Emp.	
	Ever Employed (1)	Ever Employed Last Year (2)	Fully Employed Last Year (3)	Employed in 2023m7 (4)	Log. Wage 2023m7 (5)	Last Year Log. Earnings (6)
a. Fuzzy RDD Estimate (τ_{FRD})						
Amount Collected (1000s, USD)	-0.372	0.033	2.088***	1.373*	0.010	0.028
	(0.678)	(0.531)	(0.737)	(0.721)	(0.012)	(0.021)
Robust p-value	0.667	0.941	0.007	0.057	0.182	0.181
Mean Baseline Outcome	71.80	57.14	25.84	40.79	-0.42	1.35
Effect Size (%)	-0.52%	0.06%	8.08%	3.37%		
b. Sharp RDD Estimate (τ_{SRD})						
Elig. 1st. App.	-1.217	0.107	6.885***	4.364*	0.031	0.089
	(2.219)	(1.697)	(2.361)	(2.257)	(0.037)	(0.065)
Robust p-value	0.666	0.939	0.005	0.054	0.173	0.173
Mean Baseline Outcome	71.80	57.14	25.84	40.79	-0.42	1.35
Effect Size (%)	-1.69%	0.19%	26.64%	10.70%		
c. First Stage (τ_{SRD})						
Elig. 1st. App.	3.271***	3.194***	3.298***	3.178***	3.113***	3.186***
	(0.271)	(0.181)	(0.289)	(0.244)	(0.215)	(0.254)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.47	7.35	7.53	7.52	6.21	6.83
Effect Size (%)	43.77%	43.45%	43.78%	42.26%	50.14%	46.66%
Selection of Bandwidth:						
Opt. Bandwidth	[0.017]	[0.031]	[0.016]	[0.020]	[0.042]	[0.027]
Effective Obs.	8,210	15,390	7,527	9,645	8,732	7,628

Notes: This table reports additional estimates of the effects of childhood income support on labor market outcomes for *women* in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Columns (1) through (4) correspond to variables that capture labor market attachment. Column (1) reports effects on the probability of ever being employed (i.e., ever having an employment spell lasting at least four consecutive months); column (2) on the probability of being employed for at least one month in the last year; column (3) on the probability of being employed for all 12 months in the last year; and column (4) on the probability of being employed in July 2023. All outcomes are defined as binary variables taking values 0 or 100, where 100 indicates that the corresponding condition is satisfied. Columns (5) and (6) present estimates for earnings-related variables, conditional on being employed in July 2023. Column (5) reports wages received in July 2023, while column (6) reports cumulative earnings over the period August 2022–July 2023. Both variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.7: Effects on Additional Labor Market Outcomes - Men

	Labor Market Attachment				Earnings Cond. on Emp.	
	Ever Employed (1)	Ever Employed Last Year (2)	Fully Employed Last Year (3)	Employed in 2023m7 (4)	Log. Wage 2023m7 (5)	Last Year Log. Earnings (6)
a. Fuzzy RDD Estimate (τ_{FRD})						
Amount Collected (1000s, USD)	0.169	0.944	1.002	1.872**	-0.016	0.047*
	(0.562)	(0.680)	(0.594)	(0.689)	(0.012)	(0.027)
Robust p-value	0.726	0.138	0.126	0.012	0.240	0.083
Mean Baseline Outcome	78.00	61.97	25.35	39.18	-0.25	1.48
Effect Size (%)	0.22%	1.52%	3.95%	4.78%		
b. Sharp RDD Estimate (τ_{SRD})						
Elig. 1st. App.	0.546	3.049	3.237	6.048***	-0.050	0.152*
	(1.818)	(2.182)	(1.903)	(2.176)	(0.038)	(0.083)
Robust p-value	0.722	0.131	0.118	0.009	0.229	0.073
Mean Baseline Outcome	78.00	61.97	25.35	39.18	-0.25	1.48
Effect Size (%)	0.70%	4.92%	12.77%	15.44%		
c. First Stage (τ_{SRD})						
Elig. 1st. App.	3.237***	3.228***	3.230***	3.230***	3.048***	3.230***
	(0.228)	(0.221)	(0.224)	(0.224)	(0.221)	(0.310)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.37	7.37	7.38	7.37	6.39	7.02
Effect Size (%)	43.91%	43.78%	43.77%	43.80%	47.66%	46.00%
Selection of Bandwidth:						
Opt. Bandwidth	[0.019]	[0.020]	[0.020]	[0.020]	[0.039]	[0.017]
Effective Obs.	9,864	10,379	10,166	10,129	8,292	5,229

Notes: This table reports additional estimates of the effects of childhood income support on labor market outcomes for *men* in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Columns (1) through (4) correspond to variables that capture labor market attachment. Column (1) reports effects on the probability of ever being employed (i.e., ever having an employment spell lasting at least four consecutive months); column (2) on the probability of being employed for at least one month in the last year; column (3) on the probability of being employed for all 12 months in the last year; and column (4) on the probability of being employed in July 2023. All outcomes are defined as binary variables taking values 0 or 100, where 100 indicates that the corresponding condition is satisfied. Columns (5) and (6) present estimates for earnings-related variables, conditional on being employed in July 2023. Column (5) reports wages received in July 2023, while column (6) reports cumulative earnings over the period August 2022–July 2023. Both variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table E.8: Effects on Firms' Characteristics - Women

	Log Number of Employees (1)	Log Total Wage Bill (2)	Log Median Wage (3)	Log Firm's Age Since 2005 (4)
a. Fuzzy RDD Estimate (τ_{FRD})				
Amount Collected (1000s, USD)	-0.014	-0.001	0.016	-0.010
	(0.070)	(0.079)	(0.013)	(0.021)
Robust p-value	0.671	0.791	0.205	0.591
Mean Baseline Outcome	4.85	11.60	6.66	4.68
b. Sharp RDD Estimate (τ_{SRD})				
Elig. 1st. App.	-0.043	-0.003	0.049	-0.030
	(0.211)	(0.240)	(0.040)	(0.064)
Robust p-value	0.672	0.791	0.192	0.586
Mean Baseline Outcome	4.85	11.60	6.66	4.68
c. First Stage (τ_{SRD})				
Elig. 1st. App.	3.015***	3.039***	3.118***	3.115***
	(0.322)	(0.309)	(0.255)	(0.264)
Robust p-value	0.000	0.000	0.000	0.000
Mean Baseline Outcome	6.42	6.42	6.34	6.31
Effect Size (%)	46.97%	47.36%	49.18%	49.39%
Selection of Bandwidth:				
Opt. Bandwidth	[0.022]	[0.024]	[0.032]	[0.030]
Effective Obs.	4,505	4,753	6,557	6,197

Notes: This table reports additional estimates of the effects of childhood income support on characteristics of the firms in which *women* are employed in July 2023, for those in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Column (1) reports effects on the log number of employees at the employing firm; column (2) on the log median wage paid by the employing firm; column (3) on the log of the total wage bill at the employing firm; and column (4) on the log firm age, measured as the number of consecutive months in which the employing firm has reported positive wages to the social security agency since April 2005. Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table E.9: Effects on Firms' Characteristics - Men

	Log Number of Employees (1)	Log Total Wage Bill (2)	Log Median Wage (3)	Log Firm's Age Since 2005 (4)
a. Fuzzy RDD Estimate (τ_{FRD})				
Amount Collected (1000s, USD)	-0.005	0.015	-0.014	0.048*
	(0.067)	(0.074)	(0.010)	(0.030)
Robust p-value	0.965	0.721	0.253	0.075
Mean Baseline Outcome	4.31	11.14	6.79	4.63
b. Sharp RDD Estimate (τ_{SRD})				
Elig. 1st. App.	-0.015	0.047	-0.043	0.146*
	(0.203)	(0.227)	(0.030)	(0.090)
Robust p-value	0.965	0.722	0.243	0.072
Mean Baseline Outcome	4.31	11.14	6.79	4.63
c. First Stage (τ_{SRD})				
Elig. 1st. App.	3.043***	3.044***	3.076***	3.043***
	(0.356)	(0.342)	(0.195)	(0.356)
Robust p-value	0.000	0.000	0.000	0.000
Mean Baseline Outcome	6.57	6.61	6.27	6.58
Effect Size (%)	46.35%	46.04%	49.03%	46.27%
Selection of Bandwidth:				
Opt. Bandwidth	[0.019]	[0.020]	[0.049]	[0.019]
Effective Obs.	3,775	4,021	10,511	3,778

Notes: This table reports additional estimates of the effects of childhood income support on characteristics of the firms in which *men* are employed in July 2023, for those in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Column (1) reports effects on the log number of employees at the employing firm; column (2) on the log median wage paid by the employing firm; column (3) on the log of the total wage bill at the employing firm; and column (4) on the log firm age, measured as the number of consecutive months in which the employing firm has reported positive wages to the social security agency since April 2005. Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Figure E.1: Robustness Tests: Months Worked Last Year, Women

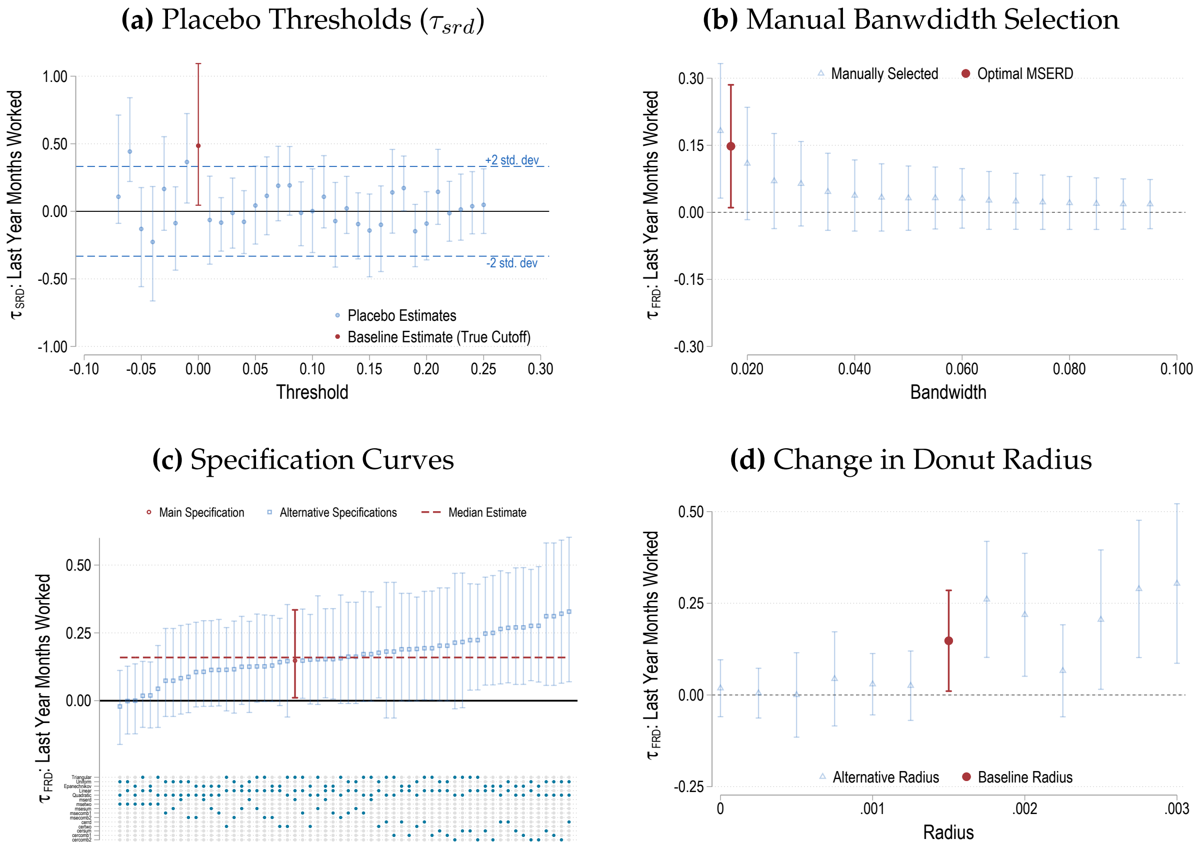

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *women's* months worked in the last year, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level.

Figure E.2: Last Year Earnings, Women

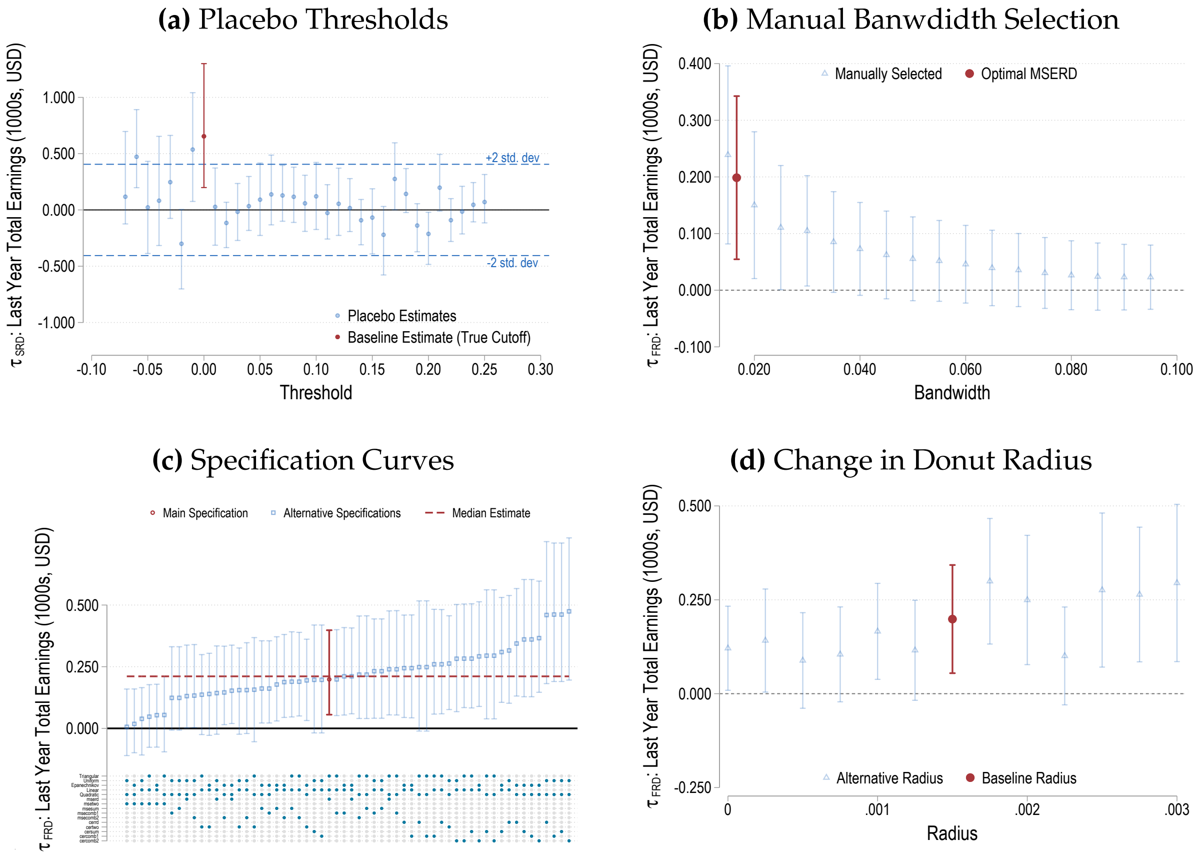

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *women's* last year earnings, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level. Earnings are expressed in 1000s of USD and are winsorized at the 95th percentile.

Figure E.3: Robustness Tests: Cumulative Months Worked, Women

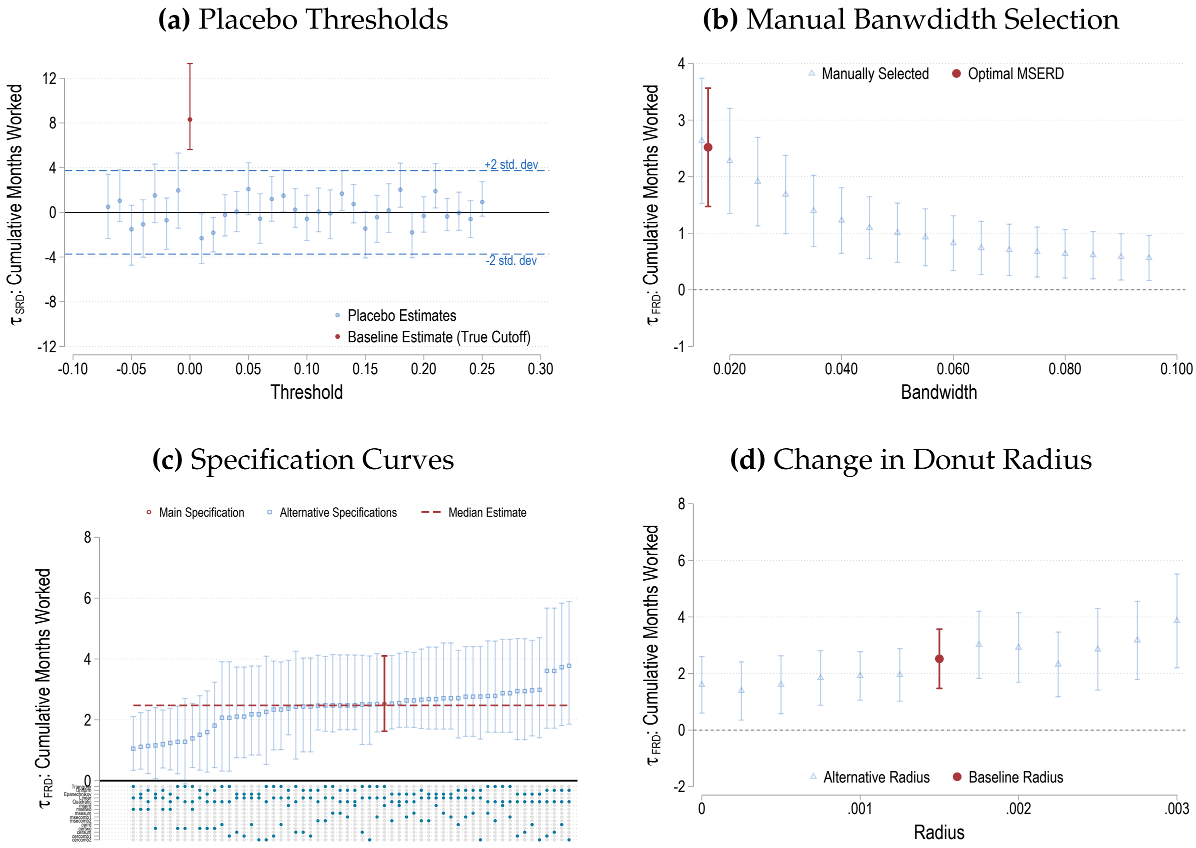

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *women's* cumulative months worked, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level.

Figure E.4: Cumulative Earnings, Women

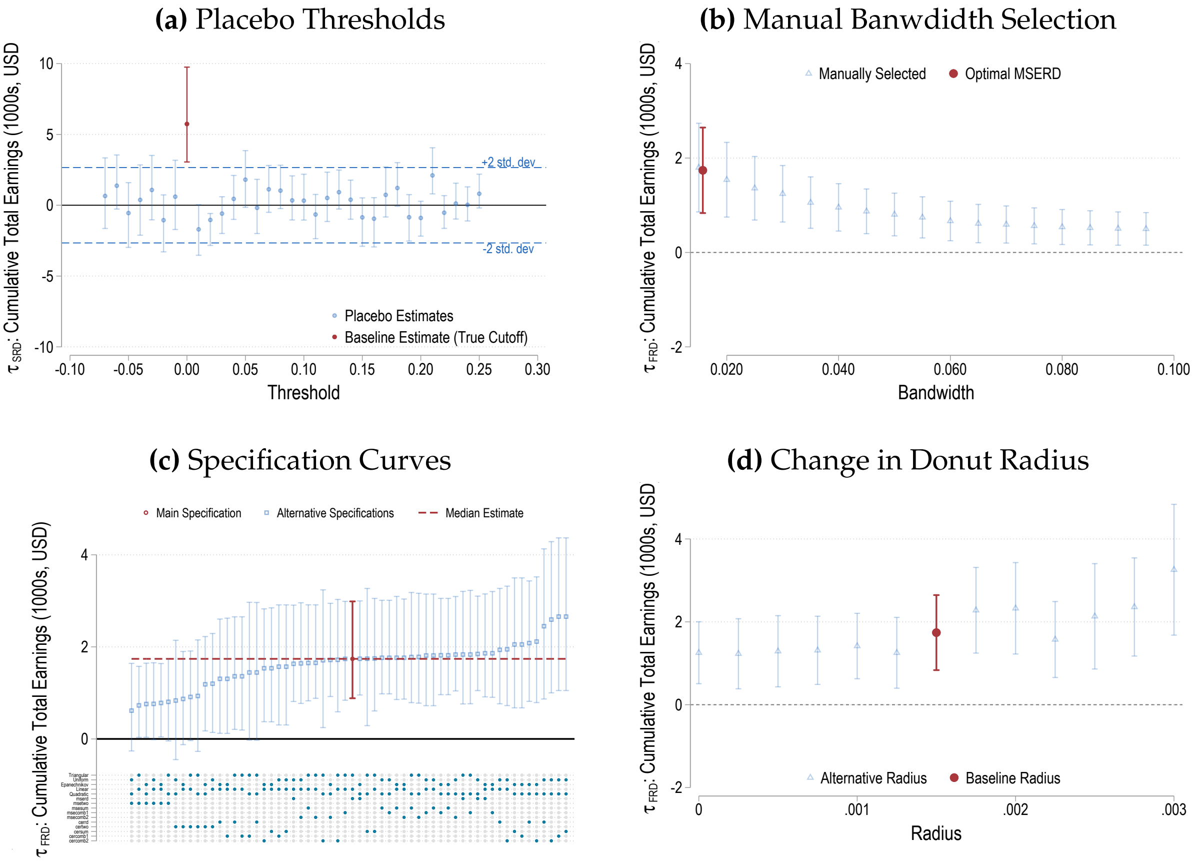

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *women's* cumulative earnings, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level. Earnings are expressed in 1000s of USD and are winsorized at the 95th percentile.

Figure E.5: Robustness Tests: Months Worked Last Year, Men

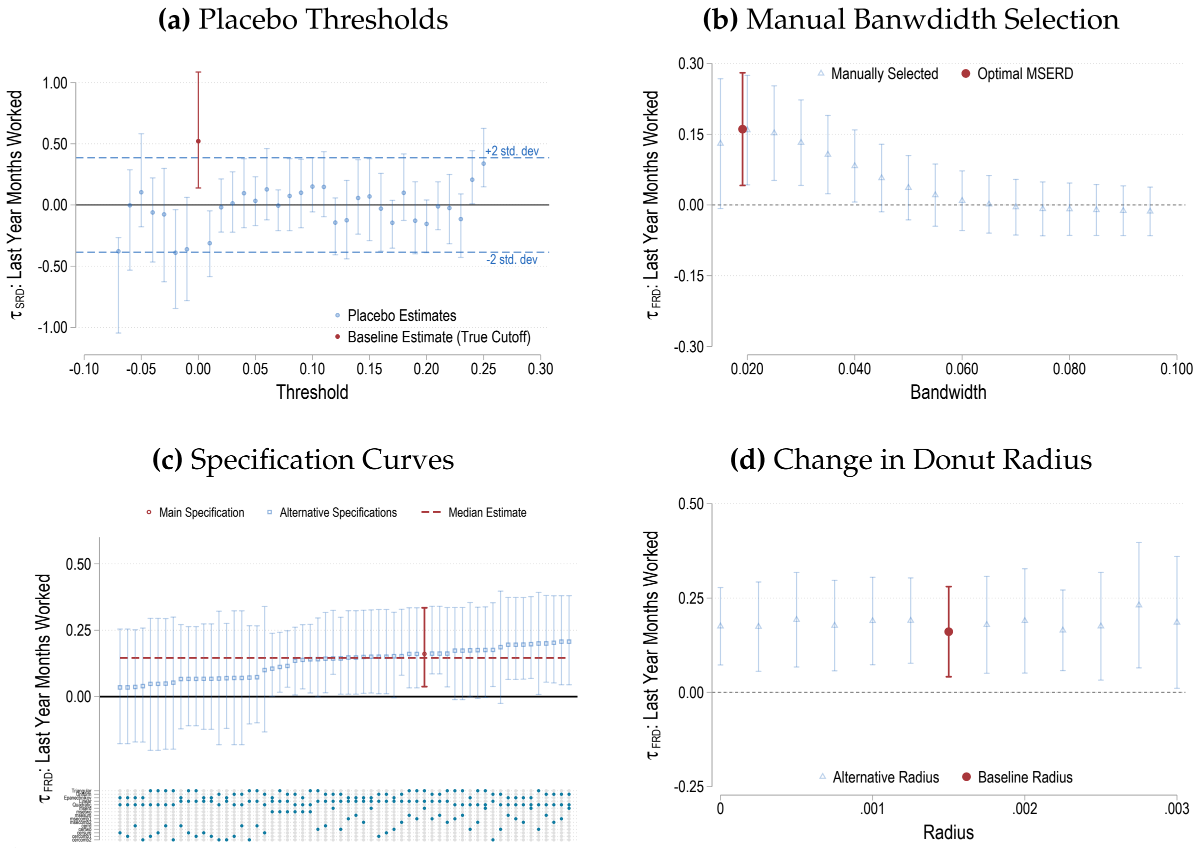

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *men's* last year months worked, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level.

Figure E.6: Last Year Earnings, Men

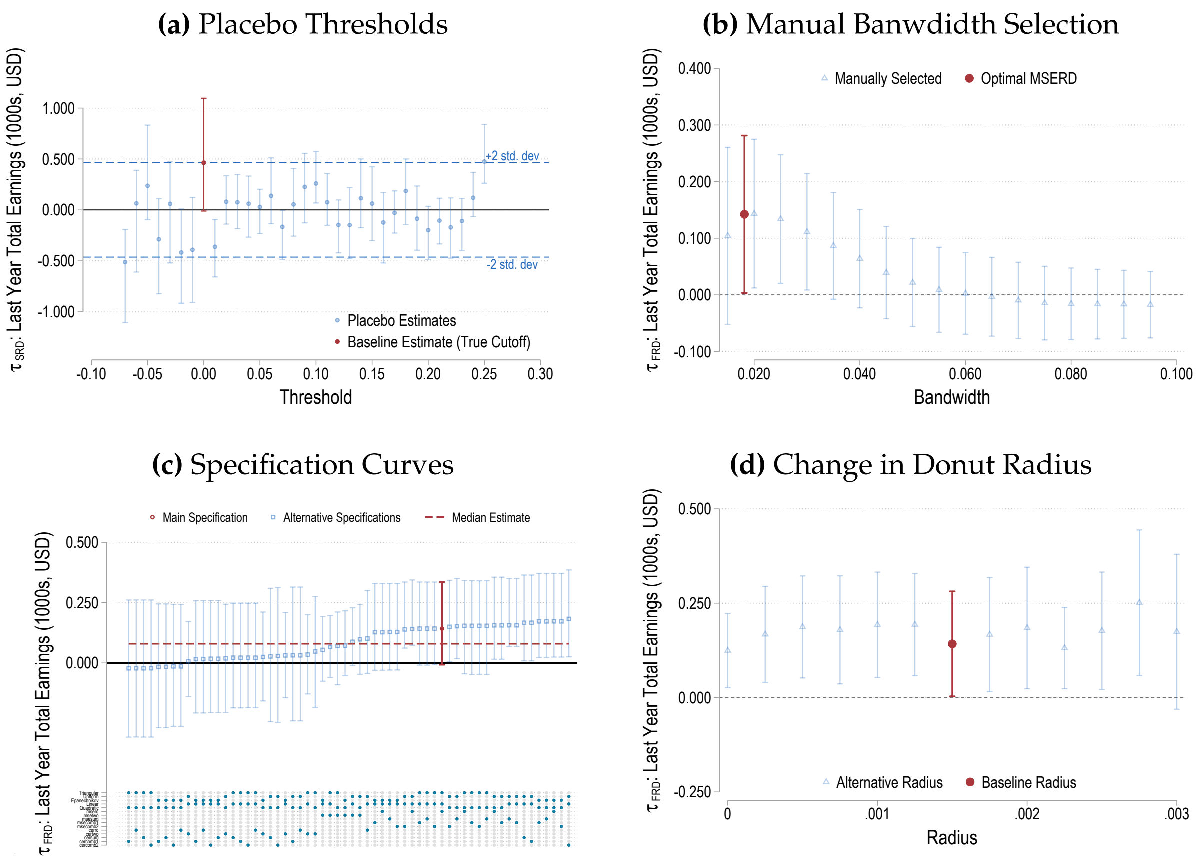

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *men's* last year earnings, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level. Earnings are expressed in 1000s of USD and are winsorized at the 95th percentile.

Figure E.7: Cumulative Months Worked, Men

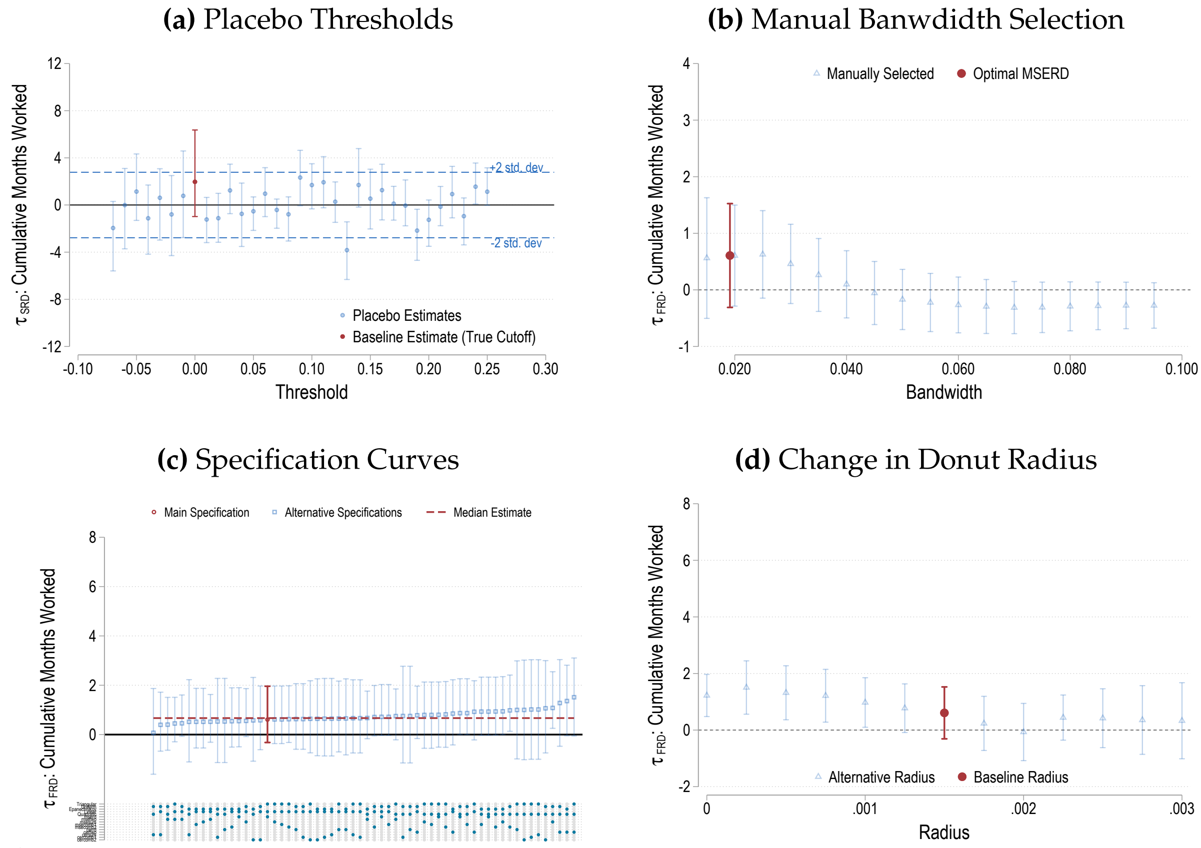

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *men's* cumulative months worked, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level.

Figure E.8: Cumulative Earnings, Men

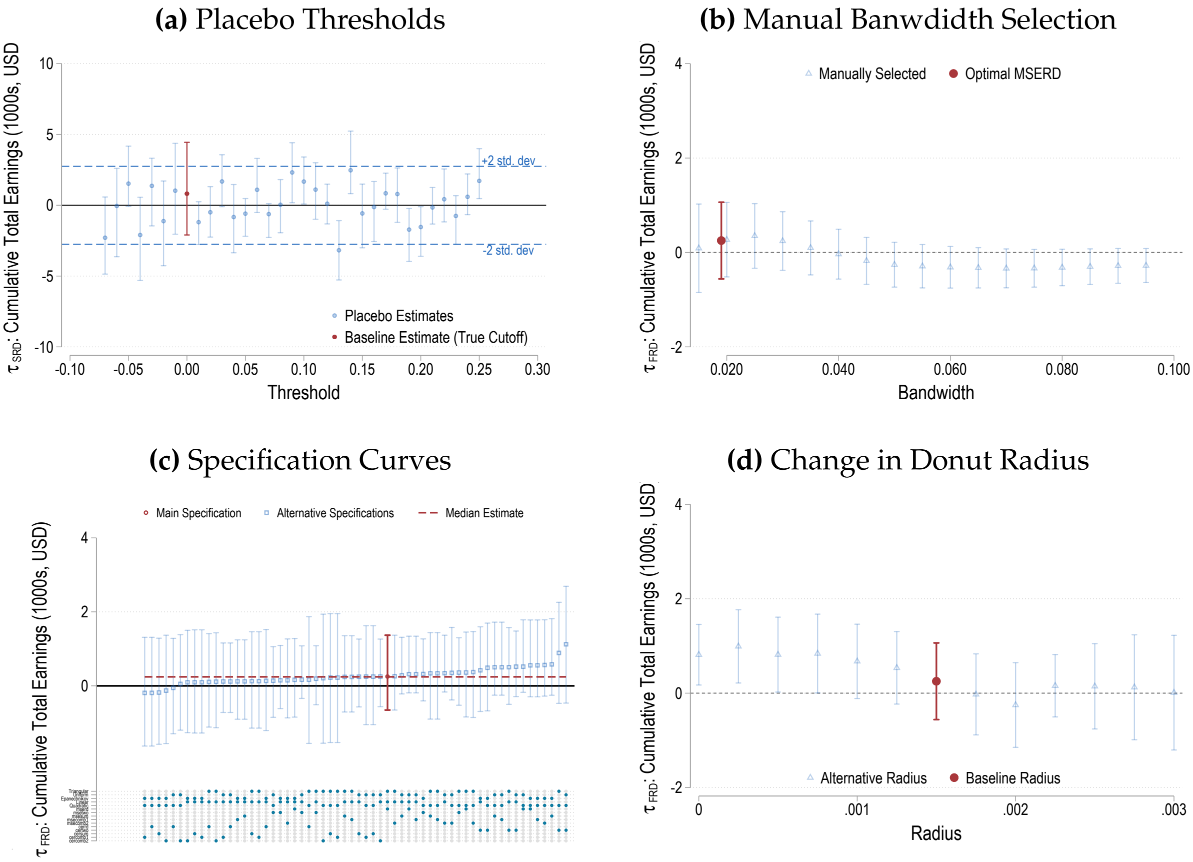

Notes: This figure reports a series of robustness checks for the estimated effects of additional childhood income support on *men's* cumulative earnings, using the *main sample*. All estimates take as the starting point the baseline model in Equations (1) or (3), and implement minimal changes to this specification to test the robustness and sensitivity of the baseline estimates to these choices. Each panel plots the different τ_{SRD} or τ_{FRD} and their corresponding confidence intervals. In all cases, the baseline estimate is highlighted separately. This baseline estimate is obtained using local linear regressions with triangular weights, bandwidths selected by minimizing the mean squared error (MSERD), is estimated at the true cutoff for Z_i^{1st} (i.e., 0), and implements a donut RDD that excludes observations within 0.0015 of the threshold. Panel (a) reports different estimates of τ_{SRD} obtained by estimating the reduced-form Equation (3) at placebo thresholds ranging from -0.07 to 0.25 in steps of 0.01, with all other specification choices held fixed. The reduced-form estimate is reported in this test because, if the RDD is valid, the first-stage at placebo thresholds should be 0, and therefore τ_{FRD} would not be valid. Dashed horizontal lines indicate plus and minus two standard deviations of the distribution of placebo estimates. Panel (b) reports different estimates of τ_{FRD} obtained by estimating Equation (1) using manually selected bandwidths ranging from 0.01 to 0.10 in steps of 0.005, with all other specification choices held fixed. Panel (c) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative specifications that vary the kernel (triangular, uniform, epanechnikov), the polynomial degree (linear, quadratic), and the bandwidth selection method (`mserd`, `msetwo`, `msesum`, `msecomb1`, `msecomb2`, `cerrd`, `certwo`, `cersum`, `cercomb1`, `cercomb2`), with all other specification choices held fixed. Each dot in the lower part of the figure indicates the specific set of estimation choices associated with the corresponding estimate, including the kernel, polynomial degree, and bandwidth selection method. Estimates are ordered along the x-axis in increasing order of τ_{FRD}. The horizontal dashed line indicates the median estimate across specifications. Panel (d) reports different estimates of τ_{FRD} obtained by estimating Equation (1) under alternative donut RDD radii ranging from 0 to 0.003 in steps of 0.00025, with all other specification choices held fixed. Confidence intervals correspond to the 90% level, and inference is based on robust standard errors clustered at the household level. Earnings are expressed in 1000s of USD and are winsorized at the 95th percentile.

Figure E.9: Effects on Employment in July 2023, by Industry - Women

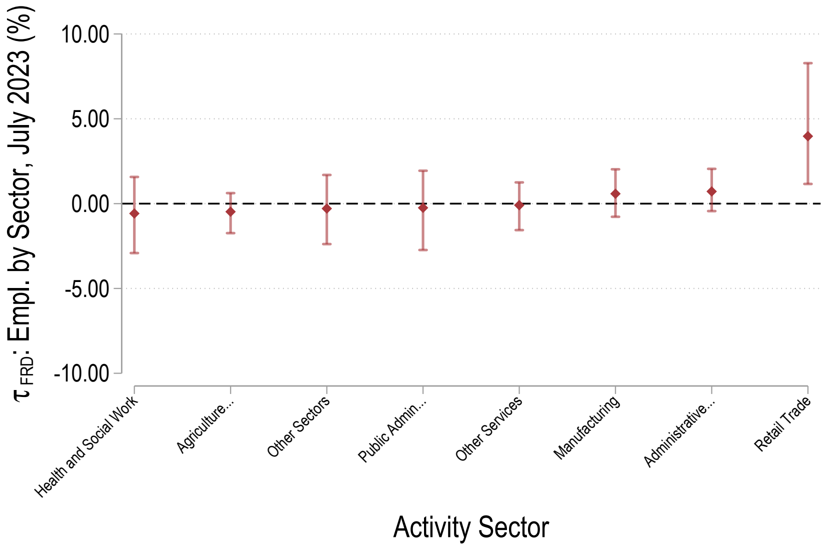

Notes: This figure reports unconditional 2SLS estimates of formal employment by activity sector for women. The vertical axis reports the 2SLS estimate τ_{FRD}, while the horizontal axis lists activity sectors. Each estimate is obtained from the baseline 2SLS specification in Equation (1) and can be interpreted as the effect of an additional USD 1,000 of childhood income support on the probability of being formally employed in sector s in July 2023 (in percentage points). Sector outcomes are constructed as binary variables equal to 100 if the individual is employed in July 2023 in sector s and equal to 0 otherwise, where the zero category includes individuals employed in a different sector as well as individuals not employed in July 2023. For women, the figure shows estimates for agriculture and extractive industries, manufacturing, retail trade, administrative and support services, public administration, human health and social work, other services, and an additional residual category ("Other sectors") that bundles sectors with fewer than 5 percent participation in the female sample (electricity and water; construction; wholesale; transportation and storage; accommodation and food; information and communication; financial and insurance; real estate; professional services; education; and arts and recreation). Estimates are based on individuals in the *main sample*, excluding observations whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold. Following Calonico et al. (2019), the optimal bandwidth (h) and bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Confidence intervals correspond to the 95 percent level, and inference is based on robust standard errors clustered at the household level.

Figure E.10: Effects on Employment in July 2023, by Industry - Men

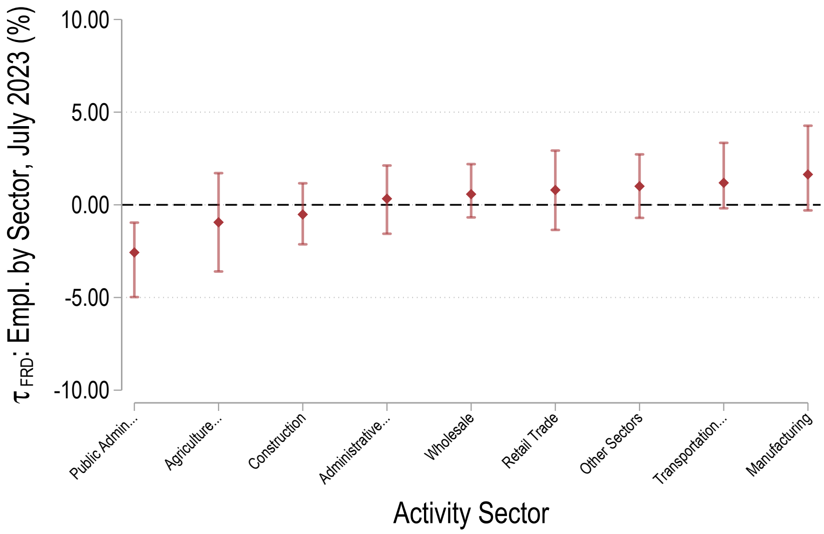

Notes: This figure reports unconditional 2SLS estimates of formal employment by activity sector for men. The vertical axis reports the 2SLS estimate τ_{FRD}, while the horizontal axis lists activity sectors. Each estimate is obtained from the baseline 2SLS specification in Equation (1) and can be interpreted as the effect of an additional USD 1,000 of childhood income support on the probability of being formally employed in sector s in July 2023 (in percentage points). Sector outcomes are constructed as binary variables equal to 100 if the individual is employed in July 2023 in sector s and equal to 0 otherwise, where the zero category includes individuals employed in a different sector as well as individuals not employed in July 2023. For men, the figure shows estimates for agriculture and extractive industries, manufacturing, construction, wholesale, retail trade, transportation and storage, administrative and support services, public administration, and an "Other sectors" category that bundles all remaining sectors with fewer than 5 percent participation in the male sample. Estimates are based on individuals in the *main sample*, excluding observations whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold. Following Calonico et al. (2019), the optimal bandwidth (h) and bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Confidence intervals correspond to the 95 percent level, and inference is based on robust standard errors clustered at the household level.

F Further Results on Mechanisms

F.1 Transition to Adulthood and Labor Market Dynamics

Effects on Other Transition-Related Outcomes. Section 4.2 reports a series of estimates capturing how additional childhood income support affected children's transitions to adulthood. Table F.1 complements those results by examining an alternative set of transition-related outcomes using the same baseline structure and specification described for Table 2. Column (1) reports estimates for the probability of becoming a mother, column (2) for the number of children, column (3) for the probability of giving birth before age 19, and column (4) for the probability of holding a relatively stable job spell before age 19.

Table F.1, column (1), shows that additional childhood income support did not affect women's likelihood of becoming a mother. Comparing women just below and just above the eligibility threshold at the time they are last observed, there are no statistically significant differences. The 2SLS estimate for the probability of becoming a mother is -1 p.p. (-1.9%) and is not statistically significant at conventional levels (p-value = 0.126). This result is relevant for two reasons. First, it indicates that additional childhood income support did not substantially affect fertility at the extensive margin over the period observed. This is consistent with effects operating primarily through changes in timing rather than through changes in the likelihood of becoming a mother. Second, the absence of large extensive-margin effects alleviates concerns that estimates on age at first birth are mostly driven by selection into motherhood. If childhood income support had significantly changed the probability of becoming a mother, then estimates on age at first birth would be a mix of compositional changes in the group of mothers with shifts in fertility timing conditional on motherhood. Column (2) complements these results with evidence on intensive-margin responses. The 2SLS estimate for the number of children is -0.032 (-3.5%), statistically significant at conventional levels (p-value = 0.031). While this could suggest a reduction in fertility at the intensive margin, the interpretation in the current sample of analysis should be more nuanced. Because women are observed, at most, at ages up to 36 and have not necessarily completed their reproductive cycles, a lower number of children at the time of observation may mechanically reflect a delayed start of motherhood rather than a reduction in completed fertility.

Table F.1, column (3), reports a substantial reduction in the probability of becoming a teen mother. The 2SLS estimate indicates that additional childhood income support decreases this probability by 2.5 p.p. (-11.6%), and is strongly statistically significant (p-value $<$ 0.001). The corresponding reduced-form discontinuity is -8.2 p.p. (-37.9%), also statistically significant at conventional levels (p-value $<$ 0.001). This result has a relatively straightforward interpretation, as the outcome is defined for all individuals in the *main sample* and is therefore not affected by compositional changes arising from selection into motherhood at later ages.

Column (4) reports estimates for a related labor-market outcome: the probability of having a relatively stable formal job before age 19. The estimated effects are also economically and statistically substantial. The 2SLS estimate indicates an increase of 2.3 p.p. (13.4%, p-value $<$ 0.001), while the corresponding reduced-form effect is 7.4 p.p. (43.6%, p-value $<$ 0.001).

Overall, the results reported in this table are consistent with sharp changes in how women transition to adulthood, characterized by delayed fertility and earlier attachment to the formal labor market. At the same time, the absence of extensive-margin fertility effects reduces concerns that the age-at-first-birth and age-at-first-job estimates are driven by selection into motherhood or employment.

Age-by-Age Estimates Section 4.2 discusses age-by-age estimates of the effects of additional childhood income support on fertility and employment outcomes, summarized in Figure 6. The figure illustrates a clear pattern: for women, fertility effects are negative and concentrated at younger ages, while employment effects are positive over roughly the same age range and attenuate at older ages. Tables F.2 and F.3 report the corresponding age-specific 2SLS and reduced-form point estimates underlying those figures for women, measured at ages 15–30 (i.e., $\gamma \in 15, 16, \ldots, 30$). In Table F.2, the coefficient reported in column (1) is obtained by estimating Equation (1) using employment at age γ as the dependent variable, restricting the sample to individuals in the *main sample* who are at least γ years old, and using the amount of cash transfers collected by the household up to age γ as the endogenous variable. Table F.3 reports the corresponding coefficients estimated from Equation (3). In each table, column (2) reports the corresponding robust standard errors, column (3) the effect size expressed as a percentage of the mean outcome for the ineligible group within the optimal bandwidth, column (4) the optimal bandwidth under the baseline strategy, and column (5) the number of effective observations. Columns (6) through (10) report analogous estimates for fertility-related outcomes, defined as having given birth by age γ. For conciseness, and given that the main purpose of these tables is to report the coefficients underlying a figure already discussed in the main text, I do not provide a detailed discussion of the results here.

Estimates Using a Balanced Sample. As discussed in Section 4.2, the baseline age-by-age estimates are constructed using age-specific samples in order to maximize the use of available information. Specifically, effects measured at age γ are estimated among individuals who are at least γ years old. As a result, the sample used to obtain each age-specific 2SLS estimate differs across ages, raising the possibility that the observed dynamic patterns reflect changes in sample composition rather than true age-specific effects. To assess this concern, Figure F.1 replicates the age-by-age analysis using a balanced sample of women who are at least 27 years old by December 2021. Panel (a) reports the corresponding 2SLS estimates for fertility and employment outcomes measured at each age using this restricted sample. The balanced sample ensures that the composition of the analysis sample remains constant across ages, although it restricts the analysis to women observed up to age 27. I choose age 27 as the cutoff because it corresponds approximately to the upper quartile of the last-observed age distribution and still provides a sufficiently large sample to implement the full RDD analysis with reasonable precision. Consistent with the main results, the dynamic patterns remain qualitatively unchanged when using the balanced sample, indicating that the age-profile effects are not driven by differences in sample composition across ages. Tables F.4 and F.5 report the corresponding point estimates, standard errors, effect sizes, optimal bandwidths, and effective numbers of observations.

For completeness, I also replicate the estimates on labor market outcomes using the balanced sample. Table F.6 shows that the results remain qualitatively and quantitatively similar when using this restricted sample. For instance, the 2SLS effect on cumulative months employed is 2.6 in the balanced sample compared to 2.5 in the *main sample*, and the effect on cumulative earnings is USD 1,932 compared to USD 1,740, respectively. Overall, the differences in magnitudes are small. Although the estimates are broadly comparable across samples, they are in general slightly lower in the balanced sample relative to the *main sample*.

Effects on Men's Transitions to Adulthood. For completeness, this appendix reports the corresponding transition-related outcomes for men. First, Figure F.2 presents the reduced-form estimates for the career-oriented transition index and its components. Second, columns (5) through (8) in Table F.1 report the corresponding point estimates for the 2SLS, reduced-form, and first stage specifications for the probability of ever having a child, the number of children, the probability of becoming a father before 19 years old, and the probability of having a job before age 19, respectively. Third, Tables F.7 and F.8 present age-by-age estimates for employment and fertility outcomes in the *main sample*. Panel (b) of Figure F.1, together with Tables F.9 and F.10, reports the corresponding results using the balanced sample. Columns (5) through (8) in Table F.6 replicate the main results on labor market outcomes for the balanced sample. Finally, Table F.11 reports estimates for additional labor market mobility and stability outcomes, analogous to those presented for women in Table 5 in the main text.

For conciseness, I do not describe each set of results in detail. Overall, the evidence is consistent with the patterns discussed for men in the main text and previous appendix sections. Across outcomes, samples, and specifications, there are no clear economically or statistically meaningful effects, except for some fertility-related outcomes, which are hard to interpret given the limitations associated with linking fathers to births described in Section 2. This reinforces the interpretation of null effects for men.

F.2 Labor Market Dynamics Around Childbirth

Point Estimates for Child Penalty by Age of First Birth Estimates. In Section 4.3, Figure 7, panel (a), summarizes the main results from the stacked DiD approach and documents heterogeneous earnings trajectories around childbirth by age at first birth. For completeness, Table F.12 reports the point estimates associated with each event time and age-at-first-birth group, together with the corresponding standard errors and the number of observations used in the estimation. Estimates are based on Equation (4) and are restricted to women who gave birth between ages 21 and 27 and whose households obtained a first application score within 0.05 of the eligibility threshold. The outcome of interest is earnings normalized relative to the within-stack average in period $t - 2$. Consistent with the graphical evidence discussed in the main text, the relative earnings penalties are larger for women who gave birth at younger ages. Moreover, penalties tend to increase over event time for early births, while they remain more stable, and in some cases attenuate, for women who gave birth at later ages.

Table F.13 reports the point estimates underlying panel (b) of Figure 7. Unlike Table F.12, which presents estimates separately for each age at first birth, this table aggregates women into two groups: those who gave birth between ages 21–23 (column (1)) and those who gave birth between ages 24–27 (column (2)). The estimation strategy follows Equation (4) and mirrors the specification used in Table F.12. The only difference is that, because stacks are pooled across age-at-first-birth groups, estimates are aggregated using weights proportional to each group's sample share, following Melentyeva and Riedel (2023) and Wing et al. (2024); Sun and Abraham (2021). Consistent with the main text, the results align with the age-specific estimates reported above, with the advantage of greater precision due to the pooling of stacks. For conciseness, I do not discuss these estimates further.

Effects on Absolute Earnings and Employment using the Stacked DiD Approach. Figure F.3 and Tables F.14 and F.15 report additional evidence using two alternative measures of labor market outcomes around childbirth. Panel (a) presents dynamic estimates using absolute earnings, rather than earnings expressed relative to pre-birth levels. Panel (b) reports analogous dynamics for employment, capturing extensive-margin labor market attachment. The estimates based on absolute earnings display a pattern that is reversed relative to the results expressed in percentage terms. In line with Melentyeva and Riedel (2023), absolute earnings losses are larger for women who gave birth at later ages, reflecting the fact that their pre-birth earnings levels are higher. When losses are expressed relative to pre-birth earnings, these baseline differences are accounted for, leading to the decreasing pattern across age at first birth documented in the main text. Importantly, the positive cumulative earnings effects of additional childhood income support reflect both higher pre-birth earnings levels and a longer period of earnings growth prior to childbirth, which more than offset the larger absolute losses at birth. Aside from this scaling difference, the qualitative dynamics remain similar, and post-birth earnings trajectories indicate that women who gave birth at later ages experience a more favorable recovery compared to women who gave birth earlier, whose earnings decline more persistently. Panel (b) complements these findings by examining employment dynamics. Although the patterns are less pronounced than for earnings, the post-birth employment trajectories suggest that part of the earnings recovery is associated with improvements in employment rates. In particular, for women who gave birth between ages 25 and 27, employment rates begin to increase as early as the year following childbirth. This pattern is consistent with stronger pre-birth labor market attachment among women who gave birth at later ages.

Point Estimates on Differential Income Dynamics by Age of First Birth and *PANES/AFAM-PE* Status. Figure 8 in Section 4.3 shows that the labor-market improvements associated with additional childhood income support are primarily driven by fertility postponement and that, once differences in age at first birth are accounted for, earnings dynamics are similar between women who received additional childhood income support and those who did not. Table F.16 reports the point estimates corresponding to the coefficients shown in the figure, based on Equation 5. Column (1) reports estimates for the reduced-form effect of *PANES/AFAM-PE* eligibility on being not yet a mother (i.e., α), column (2) reports the differential earnings dynamics relative to year $t-2$ around childbirth for ineligible mothers (i.e., γ), and column (3) reports the corresponding differential effect for eligible mothers (i.e., β). Consistent with the discussion in the main text, the estimates indicate that earnings gains for women who have not yet become mothers increase over time. However, conditional on age at first birth, eligible and ineligible mothers show very similar post-birth earnings dynamics, suggesting that differences in earnings trajectories are primarily explained by changes in fertility timing rather than differential post-birth responses by exposure to *PANES/AFAM-PE*.

Earning Dynamics Using Unstacked Data. Figure F.4 and Table report earnings and *PANES/AFAM-PE* effect dynamics around the birth event using an unstacked approach closer to Kleven et al. (2025, 2019). This unstacked DiD–RDD specification pools all mothers together and compares women at the same event time relative to birth, while allowing the reduced-form effect of program eligibility to vary flexibly over event time. Relative to the stacked specification in Equation 4, this approach does not condition on age-at-first-birth sub-events and therefore averages child penalties across all ages at first birth. This approach is useful for visualizing raw earnings patterns around childbirth and for assessing the magnitude and timing of the effects of additional childhood income support. The unstacked specification is as follows:

$$\begin{aligned} Y_{ia} = \sum_{l=-5}^{3} \Big[& \delta_l \, \mathbb{1}[a - A_i = l] \;+\; \tau_l \, D_i^{1st} \, \mathbb{1}[a - A_i = l] \\ & +\; \phi_l \, Z_i^{1st} \, \mathbb{1}[a - A_i = l] \;+\; \psi_l \, Z_i^{1st} D_i^{1st} \mathbb{1}[a - A_i = l] \Big] \\ & +\; \lambda_a \;+\; \mathbf{X_i} + \mathbf{\Lambda} + \;+\; \varepsilon_{ia}, \end{aligned} \tag{F.1}$$

where Y_{ia} denotes annual earnings for woman i at age a, expressed relative to average earnings two years before birth (i.e., normalized so that $Y_{ia} = 0$ on average at event time $l = -2$). Event time is defined as $a - A_i$, where A_i is the age at first birth. The indicator $\mathbb{1}[a - A_i = l]$ takes value one if woman i is l years away from childbirth, for $l \in \{-5, ...3\}$. D_i^{1st} is the RDD treatment indicator, equal to one for women eligible for additional childhood income support and zero otherwise, and Z_i^{1st} is the normalized running variable. The coefficients δ_l trace the evolution of earnings around childbirth for ineligible women, normalized to zero at $l = -2$. The coefficients τ_l capture, for each event time l, the differential earnings dynamics of eligible women relative to ineligible women. Hence, τ_l coefficients can be interpreted as reduced-form dynamic effects of additional childhood income support on earnings around childbirth. The interaction terms with Z_i^{1st} flexibly allow for different slopes of the running variable on either side of the threshold at each event time. Age fixed effects λ_a control for life-cycle earnings profiles common across cohorts, while $\mathbf{X_i}$ and $\mathbf{\Lambda}$ represent the control variables and fixed effects from the baseline specification. Standard errors are clustered at the individual level.

First, Figure F.4 depicts the evolution of relative earnings for ineligible women who become mothers at event time 0 (the δ_l coefficients in Equation F.1). This reflects the dynamics of labor market earnings for ineligible mothers around childbirth, in the spirit of the child-penalty literature. The patterns are as expected: women's earnings drop sharply at the time of childbirth and remain below pre-birth levels afterwards. Second, the figure depicts the differential effect for eligible mothers (the τ_l coefficients in Equation F.1). These coefficients capture the difference in earnings dynamics around birth between eligible and ineligible mothers. In other words, they represent the RDD dynamic effects of additional childhood income support around childbirth. At the same time, because the program shifts the timing of childbirth, event time $l = 0$ occurs at different ages for eligible and ineligible mothers, so these estimates combine treatment effects with changes in age at birth and should be interpreted as descriptive.

The figure shows that women who received additional income support during childhood experience increasing labor market gains up to childbirth. After childbirth, however, the effects on earnings flatten and even start to decline. Taken together, these results are consistent with the overall patterns documented in the stacked approach. The unstacked approach shows that dynamic effects of additional childhood income support clearly slow down once women become mothers. This complements the evidence from Figure 8 in the stacked specification, which shows that the gains associated with additional childhood income support are increasing while women remain childless.

F.3 Additional Evidence on Mechanism and Heterogeneity

Point Estimates for Heterogeneous Effects by Age of Exposure. In Section 4.4, Figure 9 presents heterogeneous effects of additional childhood income support by age of exposure to *PANES/AFAM-PE*. Tables F.18 through F.22 report the corresponding 2SLS, reduced-form, and first-stage point estimates underlying each panel of that figure. Specifically, Table F.18 reports the estimates for cumulative earnings, Table F.19 for cumulative months worked, Table F.20 for age at first birth, Table F.21 for age at first job, and Table F.22 for university enrollment. In all cases, the estimation strategy follows the baseline fuzzy RDD specification described in Equation (1), applied separately to groups defined by age in April 2005. For conciseness, I do not discuss these estimates further here, as the corresponding patterns were already described in the main text.

Point Estimates on Effects by Early Pregnancy and Secondary Education Level. In Section 4.4, Figure 10 presents heterogeneous effects of additional childhood income support on cumulative labor-market outcomes by early birth status and educational attainment. Table F.23 reports the corresponding 2SLS, reduced-form, and first-stage point estimates underlying panels (a) and (b) of that figure. The table implements the same baseline fuzzy RDD specification described in Equation (1), applied separately to the three groups defined in the main text: (i) women who had an early birth, (ii) women who did not have an early birth and attained a low level of education, and (iii) women who did not have an early birth and attained a high level of education. For conciseness, I do not discuss these estimates further here, as the main patterns were already described in the main text.

Figure F.1: Employment and Fertility Dynamics in the Balanced Sample

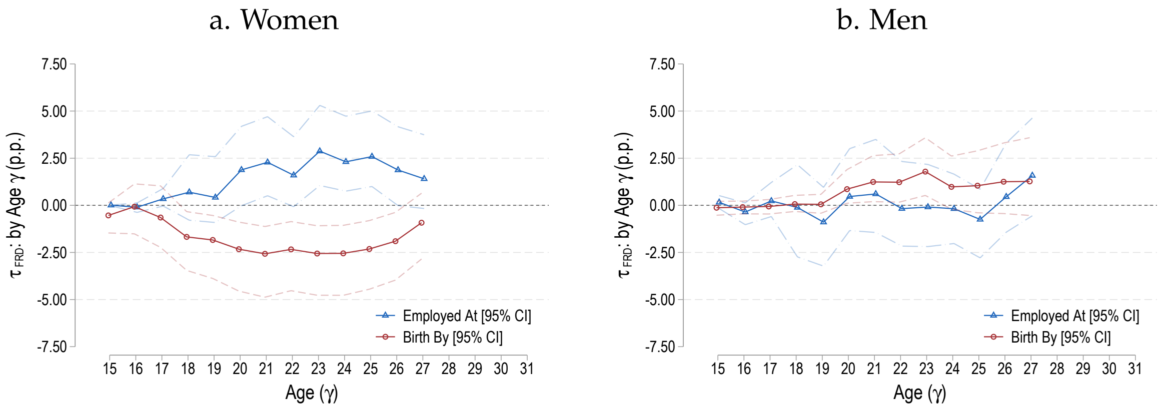

Notes: This Figure reports dynamic 2SLS estimates on employment, and fertility outcomes for the *balanced sample*. Panels (a) and (b) report dynamic 2SLS estimates of employment and fertility outcomes for women and men, respectively, as a function of age. The vertical axis reports the 2SLS estimate of the outcome measured at age γ, while the horizontal axis denotes age γ over the range 15 to 27. Each estimate corresponds to the baseline 2SLS specification in Equation (1) estimated on the *balanced sample*, , i.e., individuals who were at least 27 years old by December 2021 (see Section 4.2 for a justification of this sample), excluding individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold, and restricting the sample to individuals who are at least γ years old when last observed. For ages $\gamma < 18$, the endogenous variable is defined as the total amount of program transfers collected prior to age γ. Employment outcomes are defined as a binary variable (0, 100) that equals 100 if the individual is employed at age γ, measured as having at least four consecutive months of employment at that age, and 0 otherwise. Fertility outcomes are defined as a binary variable (0, 100) that equals 100 if the individual has given birth by age γ, and 0 otherwise. Whenever shown, confidence intervals correspond to the 95 percent level, and inference is based on robust standard errors clustered at the household level. Full estimates for Panels (a) and (b) are reported in Tables F.4 through F.10.

Figure F.2: Changes in Transition to Adulthood, Men

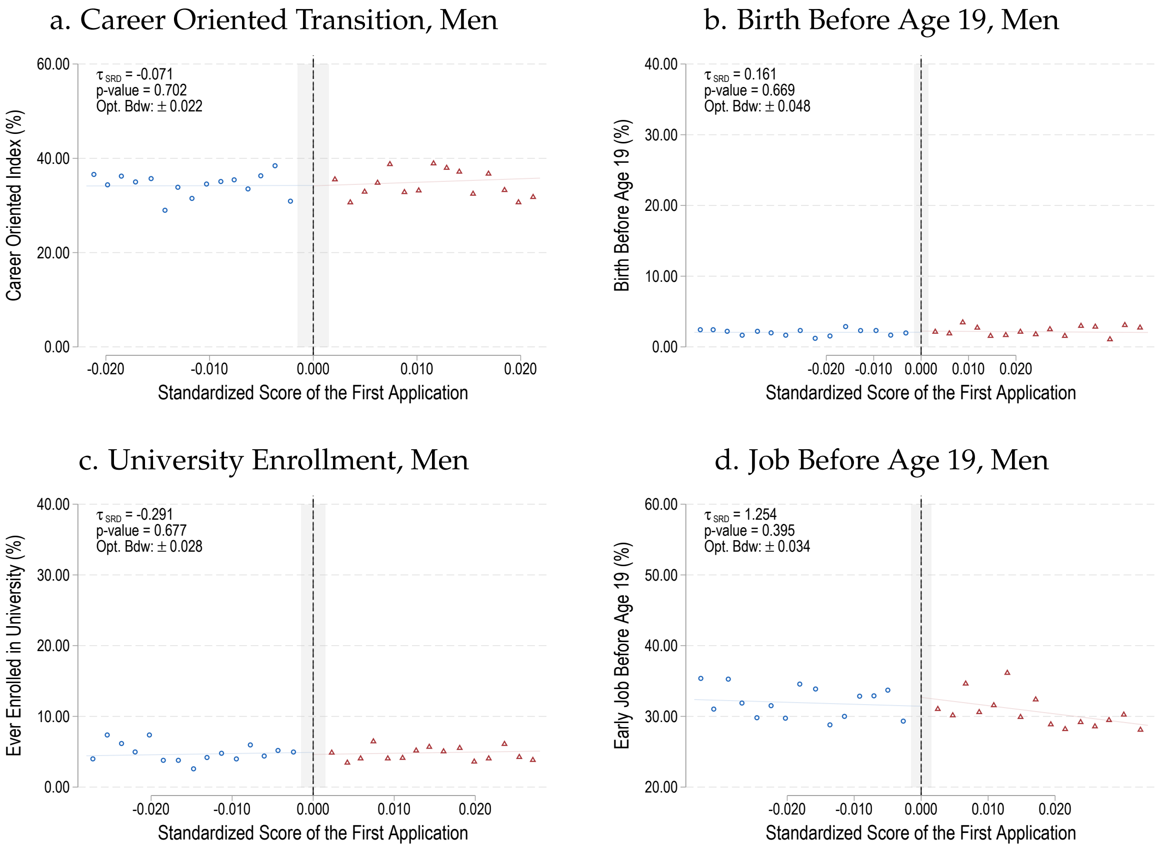

Notes: This figure reports a series of reduced-form RDD estimates for *men* in the *main sample* as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Each panel illustrates the outcome of interest as a function of the standardized poverty score at first application (Z_i^{1st}). Negative values of Z_i^{1st} indicate that the application does not meet the eligibility requirements, while positive values correspond to eligible applications. Each dot represents the bin-level average of the corresponding outcome of interest; after adjusting for the set of covariates $\mathbf{X}$ and $\mathbf{\Lambda}$ described in Section 3.1 to maintain consistency across specifications. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Outcomes are grouped into 15 quantile-spaced bins on each side of the threshold (i.e., about three times the optimal number of bins), following guidance on visual inference in RDD from Korting et al. (2023). In addition, the figure reports the point estimate of the local difference in the outcome of interest at the threshold (τ_{SRD}), based on Equation (3), together with the selected bandwidths and the continuity test p-value, shown in the upper-left corner. Each panel illustrates one of the four main outcome variables discussed in Section 4.2. Panel (a) corresponds to an index of career-oriented transition to adulthood. This index is defined as a binary variable (0, 100) that equals 100 if the individual follows a career-oriented transition to adulthood, and 0 otherwise. An individual is classified as career-oriented if two conditions hold: (i) they did not have a child before age 19, and (ii) they either held their first relatively stable job before age 19, defined as an employment spell of at least four consecutive months, or enrolled in university. Panel (b) presents analogous analysis for a binary variable (0,100) that indicates having a child before age 19, panel (c) for the probability of ever being enrolled in university as a binary variable (0, 100), and panel (d) for a binary variable (0, 100) that indicates if the individual experienced a spell of at least four consecutive months of employment before age 19. Inference is based on robust standard errors clustered at the household level. Full estimates are reported in Table 3.

Figure F.3: Child Penalty by Age of First Birth - Absolute Earnings and Employment

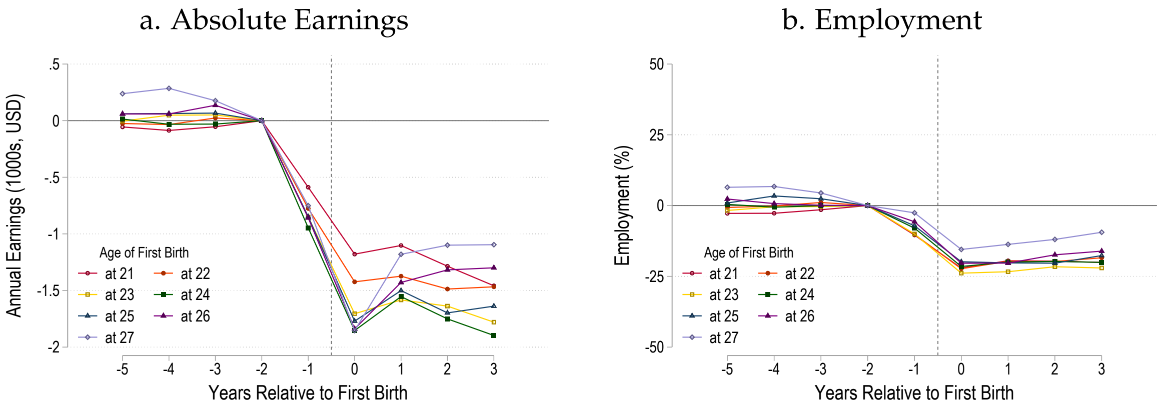

Notes: This figure reports stacked difference-in-differences estimates of labor market dynamics around childbirth for women, following the same empirical strategy described in Section 4.3. Estimates are based on Equation (4), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Each panel shows estimated coefficients $\beta_{s,l}$ as a function of years relative to first birth (l), where the horizontal axis denotes event time and the vertical axis reports the estimated effect on the outcome of interest. Panel (a) reports estimates for annual labor earnings measured in absolute terms and expressed in USD. Panel (b) reports analogous estimates for employment, defined as a binary variable (0, 100) equal to 100 if the individual is employed in a given year and 0 otherwise. Full estimates are reported in Table F.12

Figure F.4: Earnings and PANES/AFAM-PE Effects around Childbirth, Unstacked Data

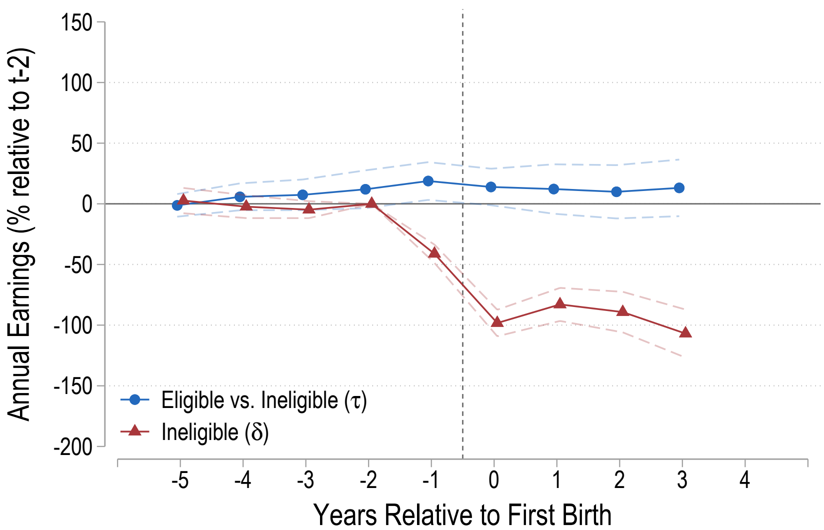

Notes: This figure reports difference-in-differences RDD estimates of earnings dynamics around childbirth for women but using unstacked data. Estimates are based on Equation (F.1), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27. Because of the RDD component of the specification, the sample is further restricted to observations within 5 p.p. of the eligibility threshold and excludes individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). The outcome of interest is annual labor earnings, expressed as a percentage of earnings two years before birth ($t = -2$), which is the omitted reference period. The figure shows two sets of event-time coefficients, δ_l and τ_l, as a function of years relative to first birth, where the horizontal axis denotes event time and the vertical axis reports the estimated effect on the outcome of interest. δ_l corresponds to the earnings dynamics of ineligible mothers around childbirth, relative to $t = -2$. τ_l corresponds to the reduced-form effect of eligibility to *PANES/AFAM-PE* around childbirth for women in the sample, based on the score obtained in the first application. Full estimates are reported in Table F.17.

Table F.1: Other Transition-Related Outcomes

	Women				Men			
	Ever Had a Child (1)	Number of Children (2)	Birth Before 19yo (3)	Job Before 19yo (4)	Ever Had a Child (5)	Number of Children (6)	Birth Before 19yo (7)	Job Before 19yo (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	-1.055	-0.032**	-2.511***	2.285***	-0.177	0.005	0.051	0.393
	(0.649)	(0.013)	(0.643)	(0.639)	(0.516)	(0.009)	(0.116)	(0.452)
Robust p-value	0.126	0.031	0.000	0.000	0.790	0.707	0.675	0.410
Mean Baseline Outcome	55.24	0.90	21.68	17.07	23.29	0.32	2.05	32.04
Effect Size (%)	-1.91%	-3.51%	-11.58%	13.39%	-0.76%	1.46%	2.47%	1.23%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	-3.373	-0.102**	-8.218***	7.448***	-0.571	0.015	0.161	1.254
	(2.038)	(0.041)	(1.915)	(1.966)	(1.661)	(0.028)	(0.367)	(1.441)
Robust p-value	0.119	0.025	0.000	0.000	0.789	0.703	0.669	0.395
Mean Baseline Outcome	55.24	0.90	21.68	17.07	23.29	0.32	2.05	32.04
Effect Size (%)	-6.11%	-11.38%	-37.90%	43.63%	-2.45%	4.75%	7.84%	3.91%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.198***	3.240***	3.273***	3.259***	3.223***	3.258***	3.169***	3.194***
	(0.250)	(0.264)	(0.276)	(0.271)	(0.214)	(0.247)	(0.131)	(0.160)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.42	7.36	7.41	7.37	7.20	7.18	6.84	7.04
Effect Size (%)	43.12%	44.04%	44.14%	44.23%	44.76%	45.40%	46.36%	45.38%
Selection of Bandwidth:								
Opt. Bandwidth	[0.019]	[0.018]	[0.017]	[0.017]	[0.022]	[0.018]	[0.048]	[0.034]
Effective Obs.	9,357	8,606	8,139	8,320	11,382	9,058	26,267	18,247

Notes: This table reports estimates of the effects of additional childhood income support on transition-related outcomes. Columns (1) through (4) correspond to women. Column (1) reports effects on the probability of ever having a child, while column (2) reports effects on the total number of births. Columns (3) and (4) present effects on the two indicators included in the career-oriented index: the probability of giving birth before age 19 and the probability of having a job (measured as being employed for at least four consecutive months) before age 19, respectively. Outcomes for columns (1), (3), and (4) are defined as binary variables taking values 0 or 100, where 100 indicates that the corresponding condition is satisfied. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.2: Employment and Fertility Dynamic Effects - 2SLS Estimates, Women

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.041	(0.042)	216.04	0.026	10,555	-0.432*	(0.284)	-14.63	0.019	7,784
16	0.195*	(0.126)	42.92	0.019	7,900	-1.047**	(0.494)	-13.01	0.018	7,649
17	0.675***	(0.206)	56.25	0.016	7,305	-1.333***	(0.567)	-9.25	0.019	8,784
18	1.618***	(0.638)	12.92	0.019	9,158	-2.510***	(0.795)	-11.60	0.017	8,147
19	1.472**	(0.804)	6.29	0.020	9,529	-1.811**	(0.859)	-6.20	0.017	8,139
20	2.052***	(0.892)	7.09	0.019	8,650	-2.044***	(0.811)	-5.58	0.022	9,982
21	2.768***	(1.002)	8.00	0.020	8,250	-1.908***	(0.786)	-4.40	0.027	11,571
22	2.657***	(1.026)	7.15	0.020	7,817	-1.533**	(0.703)	-3.14	0.032	12,950
23	2.276***	(0.763)	5.60	0.033	12,160	-1.819***	(0.761)	-3.40	0.032	11,773
24	2.893***	(1.085)	6.74	0.024	7,750	-2.605***	(1.030)	-4.48	0.024	7,829
25	2.272**	(1.084)	5.12	0.026	7,557	-2.272***	(0.979)	-3.63	0.026	7,704
26	2.145**	(1.092)	4.62	0.027	7,001	-1.496*	(0.851)	-2.28	0.032	8,278
27	1.407*	(0.998)	2.93	0.032	7,011	-0.923	(0.890)	-1.35	0.032	7,085
28	1.204	(1.003)	2.51	0.041	7,433	-0.795	(0.892)	-1.13	0.039	7,361
29	0.433	(0.739)	0.86	0.106	15,101	0.020	(0.715)	0.03	0.069	11,172
30	-0.054	(1.323)	-0.11	0.040	4,711	0.560	(1.187)	0.75	0.036	4,351

Notes: This table reports age-by-age *2SLS* estimates of the effects of additional childhood income support on *women's* employment and fertility outcomes, for ages ranging from 15 to 30. Column (0) indicates the age γ at which effects are estimated. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. The endogenous variable is additional childhood income support, measured in 1,000 USD. Separate regressions are estimated for each age γ. Column (1) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.3: Employment and Fertility Dynamic Effects - Reduced Form Estimates, Women

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.132	(0.136)	699.35	0.026	10,555	-1.373*	(0.894)	-46.48	0.019	7,784
16	0.637*	(0.406)	140.44	0.019	7,900	-3.433**	(1.576)	-42.65	0.018	7,649
17	2.336***	(0.680)	194.62	0.016	7,305	-4.500***	(1.835)	-31.22	0.019	8,784
18	5.171***	(1.975)	41.30	0.019	9,158	-8.214***	(2.368)	-37.95	0.017	8,147
19	4.683**	(2.497)	19.99	0.020	9,529	-5.927**	(2.686)	-20.31	0.017	8,139
20	6.505***	(2.763)	22.49	0.019	8,650	-6.524***	(2.487)	-17.82	0.022	9,982
21	8.428***	(2.939)	24.37	0.020	8,250	-5.895***	(2.364)	-13.59	0.027	11,571
22	8.154***	(3.080)	21.94	0.020	7,817	-4.837**	(2.189)	-9.89	0.032	12,950
23	7.063***	(2.319)	17.39	0.033	12,160	-5.661***	(2.337)	-10.57	0.032	11,773
24	8.789***	(3.223)	20.46	0.024	7,750	-8.009***	(3.097)	-13.78	0.024	7,829
25	7.062**	(3.330)	15.91	0.026	7,557	-7.133***	(3.032)	-11.39	0.026	7,704
26	6.879**	(3.440)	14.83	0.027	7,001	-4.832*	(2.742)	-7.36	0.032	8,278
27	4.440*	(3.132)	9.23	0.032	7,011	-2.938	(2.834)	-4.31	0.032	7,085
28	3.668	(3.050)	7.65	0.041	7,433	-2.427	(2.730)	-3.46	0.039	7,361
29	1.339	(2.283)	2.65	0.106	15,101	0.060	(2.189)	0.08	0.069	11,172
30	-0.163	(3.971)	-0.32	0.040	4,711	1.687	(3.575)	2.27	0.036	4,351

Notes: This table reports age-by-age *reduced-form* estimates of the effects of additional childhood income support on *women's* employment and fertility outcomes, for ages ranging from 15 to 30. Column (0) indicates the age γ at which effects are estimated. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. Separate regressions are estimated for each age γ. Column (1) reports reduced-form estimates of the sharp RDD treatment effect, τ_{SRD}, based on Equation (3) on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification reported in Table F.2. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.4: Employment and Fertility Dynamic Effects - 2SLS Estimates, Women - Balanced Sample

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.006	(0.009)	.	0.031	4,249	-0.535*	(0.406)	-15.71	0.034	4,755
16	-0.099	(0.128)	-32.25	0.034	5,647	-0.070	(0.676)	-0.82	0.033	5,607
17	0.342*	(0.236)	28.30	0.023	4,353	-0.647	(0.833)	-4.44	0.025	4,913
18	0.697	(0.886)	4.67	0.023	5,030	-1.675**	(0.796)	-7.38	0.031	6,798
19	0.418	(0.891)	1.56	0.036	7,843	-1.839***	(0.853)	-5.78	0.033	7,346
20	1.881*	(1.075)	6.17	0.028	6,242	-2.327***	(0.933)	-5.79	0.031	6,849
21	2.282**	(1.070)	6.20	0.028	6,167	-2.572***	(0.957)	-5.45	0.031	6,909
22	1.595*	(0.947)	4.05	0.034	7,516	-2.337***	(0.935)	-4.45	0.032	7,140
23	2.874***	(1.084)	6.88	0.029	6,317	-2.557***	(0.940)	-4.50	0.031	6,950
24	2.309***	(1.014)	5.15	0.030	6,521	-2.548***	(0.947)	-4.22	0.031	6,889
25	2.585***	(1.022)	5.66	0.031	6,761	-2.317***	(0.929)	-3.64	0.031	6,982
26	1.872*	(1.067)	4.04	0.029	6,434	-1.903**	(0.916)	-2.87	0.032	7,038
27	1.407*	(0.998)	2.93	0.032	7,011	-0.923	(0.890)	-1.35	0.032	7,085

Notes: This table reports age-by-age *2SLS* estimates of the effects of additional childhood income support on *women's* employment and fertility outcomes, for ages ranging from 15 to 27. Column (0) indicates the age γ at which effects are estimated. Estimates are based on women in the *balanced sample*, i.e., individuals who were at least 27 years old by December 2021 (see Section 4.2 for a justification of this sample), excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. The endogenous variable is additional childhood income support, measured in 1,000 USD. Separate regressions are estimated for each age γ. Column (1) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.5: Employment and Fertility Dynamic Effects - Reduced Form Estimates, Women - Balanced Sample

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.017	(0.026)	.	0.031	4,249	-1.595*	(1.202)	-46.84	0.034	4,755
16	-0.297	(0.384)	-96.81	0.034	5,647	-0.213	(2.045)	-2.46	0.033	5,607
17	1.096*	(0.751)	90.58	0.023	4,353	-2.088	(2.677)	-14.33	0.025	4,913
18	2.213	(2.807)	14.82	0.023	5,030	-5.347**	(2.500)	-23.55	0.031	6,798
19	1.309	(2.785)	4.88	0.036	7,843	-5.838***	(2.690)	-18.35	0.033	7,346
20	5.977**	(3.382)	19.60	0.028	6,242	-7.427***	(2.951)	-18.46	0.031	6,849
21	7.254**	(3.369)	19.71	0.028	6,167	-8.203***	(3.021)	-17.39	0.031	6,909
22	5.014*	(2.965)	12.74	0.034	7,516	-7.435***	(2.951)	-14.16	0.032	7,140
23	9.125***	(3.392)	21.83	0.029	6,317	-8.150***	(2.977)	-14.35	0.031	6,950
24	7.320***	(3.180)	16.34	0.030	6,521	-8.127***	(2.998)	-13.45	0.031	6,889
25	8.178***	(3.189)	17.90	0.031	6,761	-7.382***	(2.936)	-11.61	0.031	6,982
26	5.939**	(3.344)	12.81	0.029	6,434	-6.061**	(2.910)	-9.14	0.032	7,038
27	4.440*	(3.132)	9.23	0.032	7,011	-2.938	(2.834)	-4.31	0.032	7,085

Notes: This table reports age-by-age *reduced-form* estimates of the effects of additional childhood income support on *women's* employment and fertility outcomes, for ages ranging from 15 to 27. Column (0) indicates the age γ at which effects are estimated. Estimates are based on women in the *balanced sample*, i.e., individuals who were at least 27 years old by December 2021 (see Section 4.2 for a justification of this sample), excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. Separate regressions are estimated for each age γ. Column (1) reports reduced-form estimates of the sharp RDD treatment effect, τ_{SRD}, based on Equation (3) on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification reported in Table F.4. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.6: Labor Market Outcomes - Balanced Sample

	Women				Men			
	Last Year		Cumulative		Last Year		Cumulative	
	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings	Months Employed	Earnings
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.128	0.142	2.653***	1.932**	0.210*	0.184	0.865	0.309
	(0.093)	(0.106)	(0.904)	(0.756)	(0.118)	(0.134)	(1.065)	(0.925)
Robust p-value	0.150	0.163	0.004	0.012	0.069	0.113	0.255	0.515
Mean Baseline Outcome	5.59	4.85	58.54	39.86	5.80	5.72	75.37	54.66
Effect Size (%)	2.30%	2.92%	4.53%	4.85%	3.63%	3.22%	1.15%	0.57%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.405	0.447	8.430***	6.119***	0.624*	0.555	2.605	0.932
	(0.292)	(0.332)	(2.830)	(2.369)	(0.347)	(0.401)	(3.190)	(2.782)
Robust p-value	0.146	0.159	0.003	0.010	0.069	0.112	0.253	0.514
Mean Baseline Outcome	5.59	4.85	58.54	39.86	5.80	5.72	75.37	54.66
Effect Size (%)	7.24%	9.21%	14.40%	15.35%	10.76%	9.71%	3.46%	1.71%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.152***	3.149***	3.178***	3.167***	2.965***	3.012***	3.012***	3.012***
	(0.147)	(0.146)	(0.162)	(0.154)	(0.190)	(0.173)	(0.173)	(0.173)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	3.06	3.05	3.10	3.07	2.98	2.94	2.94	2.94
Effect Size (%)	103.19%	103.08%	102.43%	103.00%	99.46%	102.43%	102.37%	102.43%
Selection of Bandwidth:								
Opt. Bandwidth	[0.033]	[0.033]	[0.028]	[0.030]	[0.023]	[0.026]	[0.026]	[0.026]
Effective Obs.	7,165	7,312	6,221	6,668	5,146	5,934	5,942	5,958

Notes: This table reports estimates of the effects of additional childhood income support on labor market outcomes for the *balanced sample*. Columns (1) through (4) correspond to women. Column (1) reports the number of months employed in the last year observed (i.e., August 2022–July 2023), while column (2) reports total labor earnings in the last year. Columns (3) and (4) report analogous outcomes for cumulative months worked and cumulative earnings over the entire observation period. Income/earnings variables are expressed in 1000s of USD and are winsorized at the 95th percentile. Columns (5) through (8) present analogous estimates for men. Estimates are conducted separately for women and men, and are based on individuals in the *balanced sample*, i.e., individuals who were at least 27 years old by December 2021 (see Section 4.2 for a justification of this sample), excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.7: Employment and Fertility Dynamic Effects - 2SLS Estimates, Men

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.010	(0.104)	2.21	0.032	13,908	-0.001	(0.097)	-0.70	0.019	8,423
16	-0.076	(0.121)	-4.31	0.039	18,276	0.042	(0.114)	12.06	0.023	10,978
17	0.141	(0.283)	4.02	0.021	10,201	-0.058	(0.144)	-6.84	0.023	11,596
18	-0.308	(0.693)	-1.21	0.022	11,206	0.051	(0.141)	2.47	0.048	26,267
19	-0.014	(0.758)	-0.04	0.022	11,152	-0.074	(0.163)	-1.85	0.068	39,367
20	0.588	(0.832)	1.36	0.021	10,236	0.326	(0.378)	4.93	0.025	12,343
21	0.607	(0.850)	1.27	0.022	9,621	0.427	(0.486)	4.35	0.023	10,529
22	-0.215	(0.692)	-0.42	0.033	13,789	0.484	(0.454)	3.53	0.032	13,788
23	0.546	(0.726)	1.04	0.034	13,034	0.489	(0.502)	2.77	0.038	15,226
24	-0.174	(0.669)	-0.33	0.039	13,518	0.473	(0.614)	2.17	0.033	11,595
25	-0.567	(0.660)	-1.04	0.044	13,695	0.429	(0.601)	1.70	0.041	13,084
26	-0.514	(0.816)	-0.93	0.041	10,992	0.800	(0.767)	2.86	0.039	10,635
27	1.579	(1.323)	2.96	0.022	4,972	1.269	(1.052)	4.16	0.029	6,786
28	0.758	(1.216)	1.38	0.034	6,568	1.249	(1.228)	3.80	0.029	5,707
29	-0.367	(1.043)	-0.65	0.052	8,315	0.317	(0.917)	0.87	0.059	10,080
30	-0.464	(1.248)	-0.81	0.052	6,621	0.061	(0.941)	0.16	0.073	9,717

Notes: This table reports age-by-age *2SLS* estimates of the effects of additional childhood income support on *men's* employment and fertility outcomes, for ages ranging from 15 to 30. Column (0) indicates the age γ at which effects are estimated. Estimates are based on men in the *main sample*, as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. The endogenous variable is additional childhood income support, measured in 1,000 USD. Separate regressions are estimated for each age γ. Column (1) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.8: Employment and Fertility Dynamic Effects - Reduced Form Estimates, Men

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.032	(0.331)	7.04	0.032	13,908	-0.002	(0.311)	-2.26	0.019	8,423
16	-0.247	(0.393)	-14.00	0.039	18,276	0.137	(0.375)	39.76	0.023	10,978
17	0.467	(0.935)	13.32	0.021	10,201	-0.190	(0.474)	-22.59	0.023	11,596
18	-0.996	(2.242)	-3.90	0.022	11,206	0.161	(0.448)	7.84	0.048	26,267
19	-0.044	(2.450)	-0.12	0.022	11,152	-0.222	(0.487)	-5.54	0.068	39,367
20	1.915	(2.693)	4.43	0.021	10,236	1.047	(1.213)	15.85	0.025	12,343
21	1.992	(2.774)	4.17	0.022	9,621	1.397	(1.581)	14.20	0.023	10,529
22	-0.688	(2.215)	-1.36	0.033	13,789	1.543	(1.448)	11.25	0.032	13,788
23	1.740	(2.307)	3.31	0.034	13,034	1.560	(1.599)	8.84	0.038	15,226
24	-0.575	(2.213)	-1.08	0.039	13,518	1.553	(2.018)	7.12	0.033	11,595
25	-1.829	(2.133)	-3.36	0.044	13,695	1.391	(1.946)	5.52	0.041	13,084
26	-1.597	(2.541)	-2.90	0.041	10,992	2.485	(2.383)	8.87	0.039	10,635
27	4.659	(3.861)	8.73	0.022	4,972	3.840	(3.175)	12.58	0.029	6,786
28	2.174	(3.474)	3.96	0.034	6,568	3.586	(3.504)	10.92	0.029	5,707
29	-1.050	(2.988)	-1.87	0.052	8,315	0.911	(2.637)	2.50	0.059	10,080
30	-1.279	(3.435)	-2.23	0.052	6,621	0.172	(2.657)	0.45	0.073	9,717

Notes: This table reports age-by-age *reduced-form* estimates of the effects of additional childhood income support on *men's* employment and fertility outcomes, for ages ranging from 15 to 30. Column (0) indicates the age γ at which effects are estimated. Estimates are based on men in the *main sample*, as defined in Section 3, excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. Separate regressions are estimated for each age γ. Column (1) reports reduced-form estimates of the sharp RDD treatment effect, τ_{SRD}, based on Equation (3) on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification reported in Table F.7. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.9: Employment and Fertility Dynamic Effects - 2SLS Estimates, Men - Balanced Sample

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.159	(0.175)	32.77	0.053	7,844	-0.128	(0.189)	-106.30	0.034	4,874
16	-0.355	(0.295)	-15.59	0.037	6,450	-0.094	(0.172)	-37.94	0.040	7,103
17	0.232	(0.462)	5.15	0.035	7,144	-0.061	(0.195)	-11.49	0.039	8,037
18	-0.115	(1.248)	-0.38	0.023	5,184	0.064	(0.218)	4.00	0.050	11,980
19	-0.893	(1.066)	-2.05	0.031	7,163	0.053	(0.256)	1.54	0.067	16,688
20	0.466	(1.109)	0.94	0.029	6,736	0.858**	(0.456)	13.38	0.034	8,019
21	0.603	(1.260)	1.14	0.026	5,859	1.240**	(0.624)	12.53	0.030	6,995
22	-0.175	(1.145)	-0.32	0.028	6,368	1.226**	(0.654)	8.69	0.034	8,168
23	-0.090	(1.112)	-0.16	0.029	6,708	1.788***	(0.781)	9.73	0.030	7,145
24	-0.181	(0.939)	-0.33	0.039	9,158	0.979*	(0.718)	4.39	0.046	10,871
25	-0.747	(0.929)	-1.36	0.039	9,162	1.042	(0.839)	4.04	0.036	8,509
26	0.454	(1.205)	0.82	0.025	5,747	1.251	(0.955)	4.44	0.030	6,995
27	1.579	(1.323)	2.96	0.022	4,972	1.269	(1.052)	4.16	0.029	6,786

Notes: This table reports age-by-age *2SLS* estimates of the effects of additional childhood income support on *men's* employment and fertility outcomes, for ages ranging from 15 to 27. Column (0) indicates the age γ at which effects are estimated. Estimates are based on men in the *balanced sample*, i.e., individuals who were at least 27 years old by December 2021 (see Section 4.2 for a justification of this sample), excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. The endogenous variable is additional childhood income support, measured in 1,000 USD. Separate regressions are estimated for each age γ. Column (1) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.10: Employment and Fertility Dynamic Effects - Reduced Form Estimates, Men - Balanced Sample

	Dep. Var: Employed At Age					Dep. Var: Birth by Age				
Age	Estimate (1)	Std. Error (2)	Eff. Size (%) (3)	Opt. Bdw. (4)	Effective Observations (5)	Estimate (6)	Std. Error (7)	Eff. Size (%) (8)	Opt. Bdw. (9)	Effective Observations (10)
15	0.456	(0.501)	94.13	0.053	7,844	-0.368	(0.539)	-304.42	0.034	4,874
16	-1.058	(0.878)	-46.48	0.037	6,450	-0.278	(0.509)	-112.39	0.040	7,103
17	0.692	(1.372)	15.33	0.035	7,144	-0.182	(0.580)	-34.14	0.039	8,037
18	-0.340	(3.707)	-1.14	0.023	5,184	0.195	(0.660)	12.11	0.050	11,980
19	-2.709	(3.238)	-6.21	0.031	7,163	0.159	(0.765)	4.61	0.067	16,688
20	1.411	(3.347)	2.84	0.029	6,736	2.612**	(1.385)	40.75	0.034	8,019
21	1.816	(3.777)	3.43	0.026	5,859	3.754**	(1.882)	37.93	0.030	6,995
22	-0.528	(3.465)	-0.98	0.028	6,368	3.735**	(1.994)	26.47	0.034	8,168
23	-0.274	(3.366)	-0.49	0.029	6,708	5.414***	(2.353)	29.48	0.030	7,145
24	-0.549	(2.856)	-1.00	0.039	9,158	2.971*	(2.179)	13.34	0.046	10,871
25	-2.271	(2.826)	-4.12	0.039	9,162	3.172	(2.552)	12.31	0.036	8,509
26	1.363	(3.615)	2.46	0.025	5,747	3.788	(2.888)	13.43	0.030	6,995
27	4.659	(3.861)	8.73	0.022	4,972	3.840	(3.175)	12.58	0.029	6,786

Notes: This table reports age-by-age *reduced-form* estimates of the effects of additional childhood income support on *men's* employment and fertility outcomes, for ages ranging from 15 to 27. Column (0) indicates the age γ at which effects are estimated. Estimates are based on men in the *balanced sample*, i.e., individuals who were at least 27 years old by December 2021 (see Section 4.2 for a justification of this sample), excluding those whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). In addition, each row implements a further sample restriction, namely that individuals must be at least γ years old by December 2021, when fertility outcomes are last observed. Separate regressions are estimated for each age γ. Column (1) reports reduced-form estimates of the sharp RDD treatment effect, τ_{SRD}, based on Equation (3) on the outcome of interest, defined as 100 if the individual was employed for at least four consecutive months at that age, and 0 otherwise. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification reported in Table F.9. Robust standard errors clustered at the household level are reported in column (2). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively, and is based on robust p-values (Calonico et al., 2014). For reference, column (3) reports the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100. Column (4) reports the optimal bandwidth, while column (5) reports the effective number of observations used in the estimation. Columns (6) through (10) follow the same structure, but with the probability of being a parent by age γ as the outcome of interest. In this case, the variable is defined as 100 if the individual has had a child by age γ, and 0 otherwise.

Table F.11: Effects on Labor Market Mobility and Stability - Men

	Mobility History					Stability in 2023m7		
	Number of Firms (1)	Stayed Same Firm (2)	Low Mobility (3)	Moderate Mobility (4)	High Mobility (5)	Tenure in current job (6)	Log. Tenure in current job (7)	Tenure +2 Year in current job (8)
a. Fuzzy RDD Estimate (τ_{FRD})								
Amount Collected (1000s, USD)	0.024	-0.841	0.309	-0.520	0.247	-0.965	-0.024	-0.752
	(0.027)	(0.695)	(0.760)	(0.797)	(0.854)	(0.761)	(0.016)	(1.050)
Robust p-value	0.202	0.234	0.619	0.503	0.807	0.304	0.206	0.562
Mean Baseline Outcome	3.82	12.27	26.67	25.50	48.01	43.84	3.32	48.55
Effect Size (%)	0.64%	-6.85%	1.16%	-2.04%	0.51%	-2.20%		-1.55%
b. Sharp RDD Estimate (τ_{SRD})								
Elig. 1st. App.	0.077	-2.555	0.914	-1.564	0.747	-2.896	-0.073	-2.303
	(0.085)	(2.086)	(2.246)	(2.387)	(2.585)	(2.269)	(0.047)	(3.202)
Robust p-value	0.196	0.231	0.621	0.503	0.807	0.291	0.194	0.565
Mean Baseline Outcome	3.82	12.27	26.67	25.50	48.01	43.84	3.32	48.55
Effect Size (%)	2.02%	-20.82%	3.43%	-6.13%	1.56%	-6.61%		-4.74%
c. First Stage (τ_{SRD})								
Elig. 1st. App.	3.172***	3.038***	2.956***	3.009***	3.031***	3.001***	3.037***	3.061***
	(0.145)	(0.284)	(0.245)	(0.261)	(0.265)	(0.263)	(0.215)	(0.296)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	7.15	6.58	6.40	6.36	6.38	6.11	6.05	6.48
Effect Size (%)	44.37%	46.14%	46.16%	47.28%	47.53%	49.15%	50.21%	47.28%
Selection of Bandwidth:								
Opt. Bandwidth	[0.039]	[0.015]	[0.019]	[0.017]	[0.017]	[0.032]	[0.044]	[0.025]
Effective Obs.	20,760	5,867	7,736	6,999	6,825	5,577	8,059	5,030

Notes: This table reports estimates of the effects of additional childhood income support on labor market mobility and job stability outcomes for men. Columns (1) through (5) report measures of labor market mobility over the full observation period. Column (1) reports the number of different firms in which the individual held his main position. Column (2) reports the probability of remaining in the same firm when last observed. Columns (3), (4), and (5) report binary indicators of low, moderate, and high mobility, respectively. Low mobility corresponds to having one or two main employers over the entire period, moderate mobility to having three or four employers, and high mobility to having five or more. Columns (6) through (8) report measures of job stability conditional on being employed in July 2023. Column (6) reports tenure in the current job, column (7) reports log tenure in the current job, and column (8) reports a binary indicator equal to 100 if tenure in the current job is at least two years, which corresponds approximately to the median tenure in the data. Estimates are based on men in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.12: Stacked DiD - Child Penalty by Age

	Dep. Var.: Earnings Relative to t-2 (%)						
	Age 1st. Birth 21	Age 1st. Birth 22	Age 1st. Birth 23	Age 1st. Birth 24	Age 1st. Birth 25	Age 1st. Birth 26	Age 1st. Birth 27
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Event t-5	-5.396**	-1.079	-0.152	0.662	2.421	3.733	9.407**
	(2.708)	(2.857)	(2.944)	(3.158)	(3.584)	(3.782)	(4.612)
Event t-4	-8.016***	-2.343	2.417	-0.942	2.584	1.350	5.777
	(2.674)	(2.717)	(2.726)	(2.851)	(3.155)	(3.276)	(3.901)
Event t-3	-5.095**	2.068	2.848	-1.186	0.716	3.875*	4.604*
	(2.324)	(2.159)	(2.174)	(2.170)	(2.344)	(2.349)	(2.742)
Event t-1	-55.871***	-52.207***	-42.449***	-38.560***	-29.504***	-24.909***	-24.644***
	(2.433)	(2.364)	(2.218)	(2.327)	(2.370)	(2.534)	(2.899)
Event t	-112.528***	-95.549***	-86.335***	-75.196***	-59.279***	-53.375***	-51.133***
	(3.118)	(3.025)	(2.917)	(3.096)	(3.247)	(3.496)	(3.889)
Event t+1	-106.520***	-92.613***	-76.891***	-60.099***	-47.485***	-35.691***	-32.917***
	(4.350)	(4.133)	(3.993)	(4.111)	(4.343)	(4.677)	(5.295)
Event t+2	-127.362***	-101.894***	-80.692***	-68.902***	-54.244***	-35.203***	-24.010***
	(5.985)	(5.464)	(5.177)	(5.414)	(5.693)	(6.182)	(7.261)
Event t+3	-148.912***	-103.065***	-91.905***	-77.543***	-58.635***	-30.457***	-23.157
	(8.128)	(7.320)	(7.105)	(7.289)	(7.858)	(8.261)	(15.621)
Observations							
N	462,362	462,362	462,362	462,362	462,362	462,362	462,362

Notes: This table reports stacked difference-in-differences estimates of earnings dynamics around childbirth for women, following the same empirical strategy described in Section 4.3. Estimates are based on Equation (4), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Each row corresponds to an event time relative to first birth, ranging from $t = -5$ to $t = +3$. Each column reports estimates for women who give birth at a specific age, with columns corresponding to ages at first birth from 21 to 27. The dependent variable is annual labor earnings, expressed as a percentage of the within-stack average earnings two years before birth ($t = -2$), which is the omitted reference period. Reported coefficients correspond to the stacked DiD estimates $\beta_{s,l}$ shown in Panel (a) of Figure 7 in the main text. Robust standard errors clustered at the individual level are reported in parentheses. Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table F.13: Stacked DiD - Child Penalty by Age, Two Groups

	Dep. Var.: Earnings Relative to t-2 (%)	
	Age 1st. Birth 21-23 (1)	Age 1st. Birth 24-27 (2)
Event t-5	-2.457	4.522**
	(1.684)	(2.296)
Event t-4	-2.686*	2.732
	(1.626)	(2.017)
Event t-3	0.010	2.526*
	(1.357)	(1.453)
Event t-1	-52.201***	-26.062***
	(1.488)	(1.515)
Event t	-101.902***	-53.836***
	(1.977)	(2.092)
Event t+1	-95.343***	-37.812***
	(2.929)	(3.151)
Event t+2	-106.912***	-38.555***
	(4.093)	(4.427)
Event t+3	-118.249***	-40.028***
	(5.397)	(6.292)
Observations		
N	462,362	462,362

Notes: This table reports stacked difference-in-differences estimates of earnings dynamics around childbirth for women, following the same empirical strategy described in Section 4.3. Estimates are based on Equation (4), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Each row corresponds to an event time relative to first birth, ranging from $t = -5$ to $t = +3$. Columns (1) and (2) correspond to age-groups that are constructed pooling ages 21-23 and 24-27, respectively, to increase statistical precision. The dependent variable is annual labor earnings, expressed as a percentage of the within-stack average earnings two years before birth ($t = -2$), which is the omitted reference period. Reported coefficients correspond to the stacked DiD estimates $\beta_{s,l}$ shown in Panel (b) of Figure 7 in the main text. Estimates are obtained using weights constructed to recover a weighted average treatment effect across age-at-first-birth group sub-events, with weights proportional to group sample shares (Wing et al., 2024; Sun and Abraham, 2021). Robust standard errors clustered at the individual level are reported in parentheses. Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table F.14: Stacked DiD - Absolute Child Penalty by Age

	Dep. Var.: Earnings 1000s, USD						
	Age 1st. Birth 21 (1)	Age 1st. Birth 22 (2)	Age 1st. Birth 23 (3)	Age 1st. Birth 24 (4)	Age 1st. Birth 25 (5)	Age 1st. Birth 26 (6)	Age 1st. Birth 27 (7)
Event t-5	-0.056*	-0.025	-0.004	0.014	0.061	0.059	0.238
	(0.029)	(0.042)	(0.057)	(0.075)	(0.098)	(0.120)	(0.154)
Event t-4	-0.087***	-0.034	0.048	-0.032	0.063	0.059	0.285**
	(0.028)	(0.040)	(0.053)	(0.069)	(0.088)	(0.106)	(0.134)
Event t-3	-0.054**	0.024	0.050	-0.030	0.067	0.136*	0.176*
	(0.024)	(0.032)	(0.043)	(0.054)	(0.068)	(0.082)	(0.098)
Event t-1	-0.588***	-0.775***	-0.851***	-0.947***	-0.848***	-0.859***	-0.751***
	(0.026)	(0.035)	(0.044)	(0.058)	(0.071)	(0.088)	(0.107)
Event t	-1.179***	-1.424***	-1.705***	-1.854***	-1.771***	-1.839***	-1.854***
	(0.033)	(0.044)	(0.057)	(0.075)	(0.093)	(0.117)	(0.136)
Event t+1	-1.103***	-1.374***	-1.583***	-1.554***	-1.501***	-1.429***	-1.181***
	(0.046)	(0.060)	(0.077)	(0.099)	(0.123)	(0.154)	(0.191)
Event t+2	-1.285***	-1.487***	-1.638***	-1.751***	-1.698***	-1.318***	-1.100***
	(0.062)	(0.078)	(0.099)	(0.125)	(0.154)	(0.199)	(0.266)
Event t+3	-1.458***	-1.468***	-1.780***	-1.897***	-1.638***	-1.300***	-1.095*
	(0.082)	(0.102)	(0.129)	(0.160)	(0.201)	(0.270)	(0.562)
Observations							
N	462,362	462,362	462,362	462,362	462,362	462,362	462,362

Notes: This table reports stacked difference-in-differences estimates of earnings dynamics around childbirth for women, following the same empirical strategy described in Section 4.3. Estimates are based on Equation (4), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Each row corresponds to an event time relative to first birth, ranging from $t = -5$ to $t = +3$. Each column reports estimates for women who give birth at a specific age, with columns corresponding to ages at first birth from 21 to 27. The dependent variable is annual labor earnings, expressed in USD 1,000s. Reported coefficients correspond to the stacked DiD estimates $\beta_{s,l}$ shown in Panel (a) of Figure F.3. Robust standard errors clustered at the individual level are reported in parentheses. Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table F.15: Stacked DiD - Employment Child Penalty by Age

	Dep. Var.: Employed (%)						
	Age 1st. Birth 21 (1)	Age 1st. Birth 22 (2)	Age 1st. Birth 23 (3)	Age 1st. Birth 24 (4)	Age 1st. Birth 25 (5)	Age 1st. Birth 26 (6)	Age 1st. Birth 27 (7)
Event t-5	-2.761***	-0.676	-1.733	0.417	0.918	2.317	6.457**
	(0.678)	(0.850)	(1.125)	(1.450)	(1.781)	(2.052)	(2.600)
Event t-4	-2.716***	-0.170	-0.433	-0.576	3.402**	0.700	6.738***
	(0.678)	(0.889)	(1.156)	(1.374)	(1.618)	(1.880)	(2.309)
Event t-3	-1.473**	1.095	-0.345	-0.019	2.408*	0.116	4.449**
	(0.670)	(0.835)	(1.006)	(1.177)	(1.339)	(1.568)	(1.860)
Event t-1	-10.408***	-10.180***	-10.017***	-7.962***	-7.146***	-5.719***	-2.568
	(0.705)	(0.845)	(0.989)	(1.174)	(1.327)	(1.529)	(1.831)
Event t	-22.083***	-22.310***	-23.875***	-21.526***	-19.847***	-20.364***	-15.492***
	(0.775)	(0.932)	(1.132)	(1.348)	(1.561)	(1.884)	(2.167)
Event t+1	-19.487***	-19.661***	-23.366***	-19.940***	-20.257***	-20.331***	-13.706***
	(0.886)	(1.041)	(1.248)	(1.498)	(1.754)	(2.116)	(2.579)
Event t+2	-19.634***	-19.764***	-21.604***	-19.722***	-20.354***	-17.397***	-11.956***
	(0.998)	(1.181)	(1.407)	(1.698)	(1.960)	(2.455)	(3.404)
Event t+3	-20.084***	-18.448***	-22.036***	-20.074***	-17.662***	-16.101***	-9.433
	(1.176)	(1.396)	(1.672)	(1.974)	(2.397)	(3.269)	(6.900)
Observations							
N	462,362	462,362	462,362	462,362	462,362	462,362	462,362

Notes: This table reports stacked difference-in-differences estimates of earnings dynamics around childbirth for women, following the same empirical strategy described in Section 4.3. Estimates are based on Equation (4), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Each row corresponds to an event time relative to first birth, ranging from $t = -5$ to $t = +3$. Each column reports estimates for women who give birth at a specific age, with columns corresponding to ages at first birth from 21 to 27. The dependent variable is employment defined as a binary variable that takes the values 0 and 100. Reported coefficients correspond to the stacked DiD estimates $\beta_{s,l}$ shown in Panel (b) of Figure F.3. Robust standard errors clustered at the individual level are reported in parentheses. Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table F.16: Stacked DiD - RDD: Relative Child Penalty and *PANES/AFAM-PE* Dynamic Effects Around Birth

	Dep. Var.: Earnings Relative to t-2 (%)		
	Eligible Effect: Not-Yet-Mothers (α) (1)	Child Penalty: Ineligible Mothers (γ) (2)	Diff. Child Penalty: Eligible Mothers (β) (3)
Event t-5	-2.623	4.020	-2.312
	(4.649)	(4.815)	(6.776)
Event t-4	-2.668	-0.721	0.988
	(3.106)	(4.147)	(5.946)
Event t-3	4.829	11.528*	5.496
	(4.305)	(6.064)	(8.738)
Event t-1	16.648**	-28.531***	-5.926
	(7.937)	(8.318)	(11.733)
Event t	20.927**	-76.862***	-12.485
	(10.579)	(9.309)	(13.296)
Event t+1	25.549	-63.452***	-11.765
	(15.883)	(13.273)	(18.788)
Event t+2	29.012	-67.015***	-13.749
	(22.540)	(18.228)	(25.415)
Event t+3	48.884	-71.578***	-37.746
	(31.735)	(24.351)	(34.002)
Observations			
N	125,263	125,263	125,263

Notes: This table reports stacked difference-in-differences RDD estimates of earnings dynamics around childbirth for women, following the same empirical strategy described in Section 4.3. Estimates are based on Equation (5). The sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27 so that each age-at-first-birth sub-event (s) is observed over the full pre- and post-birth window. Because of the RDD component of the specification, the sample is further restricted to observations within 5 p.p. of the eligibility threshold and excludes individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach).Each row corresponds to an event time relative to first birth. The dependent variable is annual labor earnings, expressed as a percentage of the within-stack average earnings two years before birth ($t = -2$), which is the omitted reference period. Columns report estimates of the stacked DiD RDD coefficients: α_l, corresponding to the eligibility effect among not-yet-mothers; γ_l, corresponding to the child penalty for ineligible mothers; and β_l, corresponding to the differential child penalty for eligible mothers. Reported coefficients correspond to the estimates shown in Figure 8 in the main text. Estimates are obtained using weights constructed to recover a weighted average treatment effect across age-at-first-birth sub-events, with weights proportional to cohort sample shares (Wing et al., 2024; Sun and Abraham, 2021). Robust standard errors clustered at the individual level are reported in parentheses. Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table F.17: Earnings and PANES/AFAM-PE Effects around Childbirth, Unstacked Data

	Dep. Var.: Earnings Relative to t-2 (%)	
	Earnings: Ineligible Mothers (δ) (1)	Eligible Effect: Eligible Mothers (τ) (2)
Event t-5	2.662	-1.271
	(5.316)	(4.752)
Event t-4	-2.350	5.791
	(4.799)	(5.649)
Event t-3	-4.808	7.416
	(3.560)	(6.447)
Event t-1	-41.168***	18.802**
	(3.846)	(7.939)
Event t	-98.252***	13.905*
	(5.577)	(7.679)
Event t+1	-82.911***	12.200
	(6.948)	(10.396)
Event t+2	-89.225***	9.925
	(8.523)	(11.211)
Event t+3	-106.930***	13.147
	(10.094)	(11.901)
Observations		
N	62,114	62,114

Notes: This table reports difference-in-differences RDD estimates of earnings dynamics around childbirth for women using unstacked data. Estimates are based on Equation (F.1), and the sample is restricted to women in the *main sample*, as defined in Section 3, who give birth between ages 21 and 27. Because of the RDD component of the specification, the sample is further restricted to observations within 5 p.p. of the eligibility threshold and excludes individuals whose standardized poverty score at first application lies within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Each row corresponds to an event time relative to first birth. The dependent variable is annual labor earnings, expressed as a percentage of earnings two years before birth ($t = -2$), which is the omitted reference period. Column (1) reports the event-time coefficients α_l, corresponding to the earnings dynamics of ineligible mothers around childbirth, relative to $t = -2$. Column (2) reports the event-time coefficients θ_l, corresponding to the reduced-form effect of eligibility to *PANES/AFAM-PE* around childbirth for women in the sample, based on the score obtained in the first application. Reported coefficients correspond to the stacked DiD estimates shown in Panel (b) of Figure F.4. Robust standard errors clustered at the individual level are reported in parentheses. Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively.

Table F.18: Effects by Age of Exposure - Cumulative Earnings

	Age of Exposure				
	2-5 (1)	6-8 (2)	9-11 (3)	12-14 (4)	15+ (5)
a. Fuzzy RDD Estimate (τ_{FRD})					
Amount Collected (1000s, USD)	-0.422	1.173**	1.806*	0.125	2.189
	(0.220)	(0.584)	(0.781)	(0.782)	(1.613)
Robust p-value	0.115	0.043	0.053	0.661	0.153
Mean Baseline Outcome	6.44	17.10	29.13	42.51	48.06
Effect Size (%)	-6.56%	6.86%	6.20%	0.29%	4.55%
b. Sharp RDD Estimate (τ_{SRD})					
Elig. 1st. App.	-1.413*	3.298**	6.107**	0.407	6.126
	(0.701)	(1.585)	(2.579)	(2.543)	(4.493)
Robust p-value	0.080	0.031	0.042	0.661	0.153
Mean Baseline Outcome	6.44	17.10	29.13	42.51	48.06
Effect Size (%)	-21.95%	19.29%	20.96%	0.96%	12.75%
c. First Stage (τ_{SRD})					
Elig. 1st. App.	3.348***	2.812***	3.380***	3.255***	2.799***
	(0.462)	(0.362)	(0.266)	(0.138)	(0.144)
Robust p-value	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	12.82	9.80	6.09	2.33	0.53
Effect Size (%)	26.12%	28.69%	55.55%	139.60%	524.30%
Selection of Bandwidth:					
Opt. Bandwidth	[0.032]	[0.036]	[0.036]	[0.063]	[0.038]
Effective Obs.	4,185	4,007	3,914	6,051	2,138

Notes: This table reports estimates of the effects of additional childhood income support on cumulative earnings over the entire observation period for different groups of women based on age of exposure to *PANES/AFAM-PE*. Columns (1) through (5) correspond to different age groups, defined based on age in April 2005. Column (1) includes individuals aged 2–5 in April 2005, column (2) those aged 6–8, column (3) those aged 9–11, column (4) those aged 12–14, and column (5) those aged 15 or older. Cumulative earnings are expressed in 1000s of USD and are winsorized at the 95th percentile. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on cumulative earnings corresponding to estimates reported in panel (a), Figure 9, in the main text. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.19: Effects by Age of Exposure - Months Worked

	Age of Exposure				
	2-5	6-8	9-11	12-14	15+
	(1)	(2)	(3)	(4)	(5)
a. Fuzzy RDD Estimate (τ_{FRD})					
Amount Collected (1000s, USD)	-0.367	1.765**	2.078**	0.279	3.463*
	(0.290)	(0.721)	(0.920)	(0.887)	(1.906)
Robust p-value	0.277	0.014	0.040	0.580	0.055
Mean Baseline Outcome	11.02	26.30	41.88	61.32	73.27
Effect Size (%)	-3.33%	6.71%	4.96%	0.45%	4.73%
b. Sharp RDD Estimate (τ_{SRD})					
Elig. 1st. App.	-1.229	5.003***	7.047**	0.908	9.675*
	(0.955)	(1.937)	(3.034)	(2.886)	(5.282)
Robust p-value	0.236	0.008	0.031	0.581	0.054
Mean Baseline Outcome	11.02	26.30	41.88	61.32	73.27
Effect Size (%)	-11.15%	19.03%	16.83%	1.48%	13.20%
c. First Stage (τ_{SRD})					
Elig. 1st. App.	3.346***	2.834***	3.391***	3.256***	2.794***
	(0.452)	(0.379)	(0.285)	(0.137)	(0.137)
Robust p-value	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	12.87	9.79	6.14	2.32	0.50
Effect Size (%)	26.00%	28.96%	55.27%	140.29%	557.52%
Selection of Bandwidth:					
Opt. Bandwidth	[0.033]	[0.033]	[0.032]	[0.063]	[0.041]
Effective Obs.	4,335	3,692	3,459	6,106	2,329

Notes: This table reports estimates of the effects of additional childhood income support on cumulative months worked over the entire observation period for different groups of women based on age of exposure to *PANES/AFAM-PE*. Columns (1) through (5) correspond to different age groups, defined based on age in April 2005. Column (1) includes individuals aged 2–5 in April 2005, column (2) those aged 6–8, column (3) those aged 9–11, column (4) those aged 12–14, and column (5) those aged 15 or older. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on cumulative months worked corresponding to estimates reported in panel (b), Figure 9, in the main text. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.20: Effects by Age of Exposure - Age of First Birth

	Age of Exposure				
	2-5	6-8	9-11	12-14	15+
	(1)	(2)	(3)	(4)	(5)
a. Fuzzy RDD Estimate (τ_{FRD})					
Amount Collected (1000s, USD)	-0.022	0.087	0.054	0.167*	0.506**
	(0.073)	(0.109)	(0.081)	(0.088)	(0.189)
Robust p-value	0.672	0.349	0.392	0.071	0.014
Mean Baseline Outcome	17.96	19.01	19.80	20.67	22.18
Effect Size (%)	-0.12%	0.46%	0.27%	0.81%	2.28%
b. Sharp RDD Estimate (τ_{SRD})					
Elig. 1st. App.	-0.081	0.174	0.182	0.549*	1.376**
	(0.268)	(0.214)	(0.271)	(0.287)	(0.502)
Robust p-value	0.658	0.344	0.377	0.073	0.011
Mean Baseline Outcome	17.96	19.01	19.80	20.67	22.18
Effect Size (%)	-0.45%	0.91%	0.92%	2.66%	6.21%
c. First Stage (τ_{SRD})					
Elig. 1st. App.	3.673***	1.986***	3.372***	3.283***	2.721***
	(0.832)	(0.445)	(0.325)	(0.193)	(0.198)
Robust p-value	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	14.21	10.57	6.43	2.47	0.62
Effect Size (%)	25.85%	18.79%	52.47%	133.00%	441.86%
Selection of Bandwidth:					
Opt. Bandwidth	[0.031]	[0.044]	[0.037]	[0.045]	[0.028]
Effective Obs.	1,075	2,473	2,613	3,165	1,312

Notes: This table reports estimates of the effects of additional childhood income support on age at first birth for different groups of women based on age of exposure to *PANES/AFAM-PE*. Columns (1) through (5) correspond to different age groups, defined based on age in April 2005. Column (1) includes individuals aged 2–5 in April 2005, column (2) those aged 6–8, column (3) those aged 9–11, column (4) those aged 12–14, and column (5) those aged 15 or older. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on age at first birth corresponding to estimates reported in panel (c), Figure 9, in the main text. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.21: Effects by Age of Exposure - Age of First Employment

	Age of Exposure				
	2-5	6-8	9-11	12-14	15+
	(1)	(2)	(3)	(4)	(5)
a. Fuzzy RDD Estimate (τ_{FRD})					
Amount Collected (1000s, USD)	-0.060	-0.081	-0.124*	-0.187**	-0.330**
	(0.041)	(0.058)	(0.074)	(0.086)	(0.144)
Robust p-value	0.142	0.154	0.076	0.027	0.010
Mean Baseline Outcome	19.76	20.43	20.62	21.17	21.87
Effect Size (%)	-0.31%	-0.39%	-0.60%	-0.88%	-1.51%
b. Sharp RDD Estimate (τ_{SRD})					
Elig. 1st. App.	-0.204	-0.241	-0.418*	-0.593**	-0.920***
	(0.134)	(0.172)	(0.247)	(0.271)	(0.393)
Robust p-value	0.110	0.138	0.071	0.027	0.009
Mean Baseline Outcome	19.76	20.43	20.62	21.17	21.87
Effect Size (%)	-1.03%	-1.18%	-2.03%	-2.80%	-4.20%
c. First Stage (τ_{SRD})					
Elig. 1st. App.	3.373***	2.993***	3.377***	3.175***	2.783***
	(0.579)	(0.364)	(0.312)	(0.205)	(0.146)
Robust p-value	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	12.21	9.17	5.88	2.39	0.52
Effect Size (%)	27.63%	32.64%	57.40%	132.88%	538.36%
Selection of Bandwidth:					
Opt. Bandwidth	[0.039]	[0.048]	[0.033]	[0.037]	[0.038]
Effective Obs.	2,531	3,684	2,815	2,753	1,858

Notes: This table reports estimates of the effects of additional childhood income support on age at first job for different groups of women based on age of exposure to *PANES/AFAM-PE*. Columns (1) through (5) correspond to different age groups, defined based on age in April 2005. Column (1) includes individuals aged 2–5 in April 2005, column (2) those aged 6–8, column (3) those aged 9–11, column (4) those aged 12–14, and column (5) those aged 15 or older. Age at first job is defined as the age at which the individual experiences a spell of at least four consecutive months of employment and is expressed in years. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on age at first job corresponding to estimates reported in panel (c), Figure 9, in the main text. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.22: Effects by Age of Exposure - University Enrollment

	Age of Exposure				
	2-5 (1)	6-8 (2)	9-11 (3)	12-14 (4)	15+ (5)
a. Fuzzy RDD Estimate (τ_{FRD})					
Amount Collected (1000s, USD)	1.801**	0.408	0.310	0.489	-1.233
	(0.873)	(0.822)	(0.668)	(0.544)	(0.955)
Robust p-value	0.023	0.543	0.638	0.495	0.174
Mean Baseline Outcome	11.13	13.14	12.15	10.92	9.74
Effect Size (%)	16.18%	3.10%	2.55%	4.48%	-12.65%
b. Sharp RDD Estimate (τ_{SRD})					
Elig. 1st. App.	6.018**	1.114	1.061	1.593	-3.466
	(2.614)	(2.238)	(2.288)	(1.773)	(2.673)
Robust p-value	0.013	0.526	0.629	0.498	0.172
Mean Baseline Outcome	11.13	13.14	12.15	10.92	9.74
Effect Size (%)	54.06%	8.48%	8.74%	14.58%	-35.57%
c. First Stage (τ_{SRD})					
Elig. 1st. App.	3.341***	2.733***	3.426***	3.259***	2.812***
	(0.573)	(0.342)	(0.261)	(0.136)	(0.140)
Robust p-value	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	12.65	9.62	6.01	2.30	0.54
Effect Size (%)	26.42%	28.40%	56.97%	141.50%	523.85%
Selection of Bandwidth:					
Opt. Bandwidth	[0.023]	[0.040]	[0.037]	[0.063]	[0.039]
Effective Obs.	3,040	4,445	4,031	6,208	2,237

Notes: This table reports estimates of the effects of additional childhood income support on the probability of ever being enrolled in university for different groups of women based on age of exposure to *PANES/AFAM-PE*. Columns (1) through (5) correspond to different age groups, defined based on age in April 2005. Column (1) includes individuals aged 2–5 in April 2005, column (2) those aged 6–8, column (3) those aged 9–11, column (4) those aged 12–14, and column (5) those aged 15 or older. Note that, because university enrollment data are available only through 2020, estimates for the youngest cohort mostly reflect responses among children who were 3–5 years old when *PANES/AFAM-PE* was rolled out, as discussed in Section 2. University enrollment is defined as a binary variable (0, 100) that equals 100 if the woman was ever enrolled in university, and 0 otherwise. Estimates are based on women in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the probability of ever being enrolled in university corresponding to estimates reported in panel (d), Figure 9, in the main text. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

Table F.23: Labor Market Effects by Early Pregnancy and Education Level

	Dep. Var.: Cumulative Total Earnings (1000s, USD)			Dep. Var.: Cumulative: Months Worked		
	Early Birth (1)	No Early Birth Max Educ.: Middle School (2)	No Early Birth Max. Educ: High Educ. (3)	Early Birth (4)	No Early Birth Max Educ.: Middle School (5)	No Early Birth Max. Educ: High Educ. (6)
a. Fuzzy RDD Estimate (τ_{FRD})						
Amount Collected (1000s, USD)	0.224	1.240**	1.635**	0.726	1.778**	2.568***
	(0.570)	(0.585)	(0.695)	(0.732)	(0.775)	(0.781)
Robust p-value	0.580	0.042	0.014	0.288	0.023	0.001
Mean Baseline Outcome	16.20	19.01	33.00	27.01	32.11	46.16
Effect Size (%)	1.38%	6.52%	4.95%	2.69%	5.54%	5.56%
b. Sharp RDD Estimate (τ_{SRD})						
Elig. 1st. App.	0.667	3.576**	5.694**	2.169	5.128**	9.267***
	(1.697)	(1.675)	(2.358)	(2.172)	(2.204)	(2.647)
Robust p-value	0.568	0.038	0.012	0.266	0.020	0.000
Mean Baseline Outcome	16.20	19.01	33.00	27.01	32.11	46.16
Effect Size (%)	4.12%	18.81%	17.26%	8.03%	15.97%	20.08%
c. First Stage (τ_{SRD})						
Elig. 1st. App.	2.980***	2.883***	3.483***	2.987***	2.884***	3.609***
	(0.309)	(0.267)	(0.383)	(0.311)	(0.267)	(0.413)
Robust p-value	0.000	0.000	0.000	0.000	0.000	0.000
Mean Baseline Outcome	8.40	7.00	7.14	8.40	7.00	7.10
Effect Size (%)	35.49%	41.19%	48.78%	35.58%	41.19%	50.80%
Selection of Bandwidth:						
Opt. Bandwidth	[0.039]	[0.036]	[0.019]	[0.039]	[0.036]	[0.017]
Effective Obs.	4,107	5,755	4,191	4,060	5,725	3,741

Notes: This table reports estimates of the effects of additional childhood income support on cumulative earnings and cumulative months worked over the entire observation period for different groups of women defined by fertility and educational trajectories. Columns (1) through (3) report estimates for cumulative earnings, while columns (4) through (6) report estimates for cumulative months worked. Within each outcome, columns correspond to mutually exclusive groups defined as follows: column (1) and column (4) include women who experienced an early birth, defined as having a first child at age 19 or earlier; column (2) and column (5) include women with no early birth whose maximum secondary education grade corresponds to enrollment in middle school; and column (3) and column (6) include women with no early birth whose maximum secondary education grade corresponds to enrollment in high school. Cumulative earnings are expressed in 1000s of USD and are winsorized at the 95th percentile. Estimates are based on individuals in the *main sample*, as defined in Section 3, excluding those with a standardized poverty score at first application within 0.0015 of the eligibility threshold (see Section 3.1 for a justification of this donut RDD approach). Panel (a) reports 2SLS estimates of the fuzzy RDD treatment effect, τ_{FRD}, based on Equation (1), interpreted as the effect of an additional USD 1,000 of childhood income support on the corresponding outcome of interest. Panel (b) reports reduced-form estimates, τ_{SRD}, based on Equation (3). Panel (c) reports first-stage estimates, δ, based on Equation (2). Following Calonico et al. (2019), the optimal bandwidth (h) and the bias bandwidth (b) are selected by minimizing the mean squared error (MSERD) using the default options in `rdrobust` and implementing local linear regressions with triangular weights. To maintain a consistent sample across the 2SLS, reduced-form, and first-stage estimations, the optimal and bias bandwidths are computed for the fuzzy RDD specification in Equation (1) and then held fixed for the reduced-form and first-stage estimates. Both the optimal bandwidth used across panels, as well as the effective number of observations used in the estimation, are reported in the bottom rows of the table. Robust standard errors clustered at the household level are reported in parentheses, together with the robust p-value, in each panel of the table (Calonico et al., 2014). Statistical significance is indicated by ***, **, and *, corresponding to the 1%, 5%, and 10% levels, respectively. For reference, each panel also reports the mean baseline outcome, computed as the mean of the corresponding outcome of interest for *ineligible* individuals within the optimal bandwidth range, as well as the effect size expressed in percent terms, computed as the ratio of the corresponding coefficient of interest τ to the mean baseline outcome, multiplied by 100.

G Further Details on the Conceptual Framework

Income support during childhood, including cash transfer programs, can affect household decisions and children's outcomes through standard income effects and, when benefits are tied to specific behaviors, through substitution effects. In practice, these mechanisms operate through changes in financial constraints, expectations, parental investments, early-life education and fertility decisions, social environment, and parental stress, which interact dynamically over time and are particularly relevant in the context studied here (Page, 2024). To structure the main ideas, this section presents a simple model of early-life income, career investment, and fertility timing. The purpose of the model is to clarify how additional income during childhood shapes incentives during the transition to adulthood and to guide the interpretation of the empirical results. While intentionally stylized, the framework is flexible enough to accommodate more specific mechanisms.

The framework builds on the career–fertility timing model in Doepke et al. (2023), which formalizes the idea that the opportunity cost of having children at young ages includes not only foregone wages, but also missed opportunities for career investment. In that framework, women may delay childbearing to invest in their careers, but doing so is risky because fertility at later ages is uncertain and increasingly costly.[60] The model predicts sorting into three fertility paths: early motherhood, delayed motherhood, and childlessness.

The basic structure of the model remains unchanged, but I extend it by introducing a childhood period that determines initial conditions and restrictions under which women transition to adulthood. In a nutshell, childhood income affects adulthood decisions through two channels: human capital accumulation, which influences entry wages and incentives for career investment, and exposure to very early fertility, which can constrain subsequent fertility timing and career investment choices. The structure of adulthood decisions is otherwise unchanged.

Period 0: Childhood income, schooling, and early fertility risk

Period 0 represents childhood and teenage years, during which decisions are made within the household. This period is introduced to make explicit how early-life income shapes the initial conditions women face at the beginning of adulthood. Childhood income affects adulthood outcomes through two channels: schooling investment and exposure to early fertility risk. For simplicity, period 0 only determines the initial conditions at the start of period 1, when individuals become independent, and I therefore abstract from a full intergenerational optimization problem.

Households allocate resources between consumption c_0 and schooling investment s_0. Preferences are given by

$$u_0(c_0, s_0) = \log(c_0) + \delta \log(s_0), \tag{G.1}$$

where $\delta > 0$ implies that schooling is a normal good. Modeling schooling as a normal good is consistent with a large literature on education and human capital investment (e.g., Todd and Wolpin 2006, 2008; Keane and Wolpin 2010).[61] Schooling has a monetary cost p, which increases

[60] For instance, fertility decisions in adulthood are influenced by rising biological and psychological costs, as well as higher pregnancy risks at later ages (Schmidt et al., 2012; Gustafsson, 2001)

[61] It is beyond the scope of this paper to discuss the non-pecuniary benefits of schooling or whether it should be

by a factor $(1+\varphi)$, with $\varphi > 0$ if a period-0 fertility shock occurs. For simplicity, I refer to births in period 0 as childhood or teenage births. Conceptually, they represent births that happened before women's transition to adulthood. Let $n_0 \in \{0,1\}$ denote whether a *teenage* birth occurs in period 0. The household budget constraint is therefore

$$c_0 + p(1+\varphi n_0)s_0 = y, \tag{G.2}$$

where y denotes household income. For our population of interest, i.e., those who were children in period 0, y is what I refer to as childhood income. The parameter φ captures tighter time and resource constraints associated with early childbearing, such as additional childcare expenditures, transportation or supervision costs related to attending school, or foregone earnings of other household members providing childcare support (Klepinger et al., 1999).[62]

Teenage fertility is not a choice variable. Instead, it occurs with probability

$$\Pr(n_0 = 1 \mid y) = r(y), \qquad r'(y) < 0, \tag{G.3}$$

so higher household income reduces exposure to period 0 fertility risk. This formulation summarizes a range of mechanisms documented in the literature. Higher-income households may face fewer financial barriers to contraceptive use and reproductive health services (Bailey, 2006; Kearney and Levine, 2009; Bailey et al., 2023), provide greater parental supervision and safer environments, or raise the perceived returns to schooling and future labor market participation, increasing the opportunity cost of an unintended early pregnancy (Black et al., 2008). More broadly, it is consistent with evidence that economic conditions and future opportunities contribute to shaping teenage fertility beyond preferences alone (Kearney and Levine, 2012, 2014). This formulation is intentionally reduced-form: rather than modeling parental choices over these mechanisms explicitly, the function $r(y)$ summarizes how childhood economic conditions translate into exposure to teenage childbearing.

Period 0 determines schooling investment and the realization of teenage fertility, both of which enter period 1 as predetermined initial conditions. The household problem can be written as

$$\max_{s_0>0} \log\Big(y - p(1+\varphi n_0)s_0\Big) + \delta \log(s_0). \tag{G.4}$$

The first-order condition is

$$-\frac{p(1+\varphi n_0)}{y - p(1+\varphi n_0)s_0} + \frac{\delta}{s_0} = 0. \tag{G.5}$$

Solving (G.5) yields the optimal schooling choice

$$s_0^*(y, n_0) = \frac{\delta}{1+\delta} \cdot \frac{y}{p(1+\varphi n_0)}. \tag{G.6}$$

considered a (normal) consumption good. Oreopoulos (2011) and MacLeod and Urquiola (2019) provide in-depth reviews of this discussion.

[62] An alternative formulation would be to include a direct disutility of schooling effort in the utility function. However, as emphasized by Keane and Wolpin (2010), monetary costs associated with particular choices are generally not distinguishable from psychic or effort costs when they are unobserved, making the distinction largely arbitrary. I therefore model the effects of early fertility through the budget constraint, which provides a transparent representation of how early-life shocks raise the marginal cost of schooling without affecting the qualitative implications of the model.

Thus, conditional on the realization of teenage fertility, schooling investment is increasing in household income and decreasing in the effective cost of schooling, which is higher if a teen birth happens. In particular,

$$s_0^*(y,1) = \frac{1}{1+\varphi}\, s_0^*(y,0). \tag{G.7}$$

Periods 1 and 2: career investment and fertility timing

Period 1 can be thought of as the *transition to adulthood*. As mentioned above, the main assumptions and modeling choices follow Doepke et al. (2023). However, the introduction of period 0 decisions and shocks modifies the incentives and constraints under which women choose fertility timing and career investment through their effects on initial conditions at the start of adulthood in two ways.

First, childhood schooling affects labor market returns both through entry wages and through wages later in life. At the beginning of period 1, women are characterized by realized schooling s_0, determined in period 0 and therefore given by the optimal rule: $s_0 \equiv s_0^*(y, n_0)$. Women enter period 1 with a baseline wage that depends on this realized level of schooling. Childhood conditions therefore affect the opportunity cost of motherhood and career investment:

$$w_1 = w_1(s_0), \qquad w_1'(s_0) > 0. \tag{G.8}$$

Second, teenage fertility realizations constrain subsequent fertility timing decisions. If $n_0 = 1$, the woman enters period 1 having already had a birth before adulthood, so fertility timing is no longer a choice. In this case, I set $n_1 = 1$, meaning that the period-1 time cost of childrearing applies to women who had a teenage birth before adulthood. If $n_0 = 0$, the woman can choose whether and when to have a child. As in Doepke et al. (2023), women can have at most one child; therefore, conditional on ever having a child, if $n_1 = 1$ then $n_2 = 0$, and vice versa. Mothers derive the same utility from a child regardless of timing.

Women allocate time between work, children, and career investment. Let $e_1 \in [0,1]$ denote career investment in period 1, and let $n_1 \in \{0,1\}$ indicate whether a birth occurs in period 1. Having a child requires a fraction $\phi \in (0,1)$ of available time, so period 1 consumption is

$$c_1 = w_1(s_0)\big(1 - e_1 - \phi n_1\big). \tag{G.9}$$

Career investment e_1 should be interpreted broadly as any use of time during the transition to adulthood that raises future wages but is not fully remunerated at the current stage. This includes formal training and education, on-the-job learning, job search for better matches, or accepting lower current earnings in exchange for future career opportunities (Goldin, 2014).

Period 2 corresponds to adulthood. Career investment undertaken in period 1 raises wages at this stage of life through a human-capital accumulation technology that follows Doepke et al. (2023), with one modification: both the scale and the curvature of returns to career investment depend on realized childhood schooling,

$$w_2 = w_1(s_0)\,\kappa(s_0)\,e_1^{\gamma(s_0)}, \qquad \kappa'(s_0) > 0, \quad \gamma'(s_0) > 0. \tag{G.10}$$

Here, $\kappa(s_0)$ captures the productivity scale of career investments, while $\gamma(s_0)$ captures the elasticity of period-2 wages with respect to career investment. Since $e_1 \in (0,1)$, increases in $\gamma(s_0)$ mechanically reduce $e_1^{\gamma(s_0)}$ holding career investment fixed. I therefore assume that the increase in the productivity scale $\kappa(s_0)$ is sufficiently strong to dominate this mechanical offset, so that higher schooling raises period-2 wages.

Thus, childhood schooling increases both entry wages and the returns to career investment during early adulthood, introducing a dynamic complementarity between early-life human capital and later career investments (Heckman, 2006; Cunha and Heckman, 2007; Cunha et al., 2010). Intuitively, additional childhood schooling may give access to career paths where investments have higher payoffs and steeper returns. For instance, some investments during the transition to adulthood, such as tertiary education, require substantial time commitments. These investments are only worthwhile for individuals who enter adulthood with a sufficiently high s_0. In the model, this complementarity is captured by assuming that higher s_0 raises both the productivity scale of career investment, $\kappa(s_0)$, and the elasticity of period-2 wages with respect to that investment, $\gamma(s_0)$.

Women may choose to attempt a birth in period 2 ($n_2 \in \{0,1\}$) rather than in period 1. If $n_2 = 1$, a birth occurs with probability $\pi \in (0,1)$; with probability $1-\pi$, the woman remains childless due to fertility risk. When a period-2 birth is realized, consumption is scaled by the time cost of childbearing:

$$c_{2,1} = w_1(s_0)\,\kappa(s_0)\,e_1^{\gamma(s_0)}(1-\phi), \tag{G.11}$$

whereas if no period-2 birth is realized,

$$c_{2,0} = w_1(s_0)\,\kappa(s_0)\,e_1^{\gamma(s_0)}. \tag{G.12}$$

Preferences are logarithmic in consumption in periods 1 and 2 and quasi-linear in children. A child yields utility v, which varies across women and is independent of the timing of birth. Because period 2 fertility is risky, I write expected lifetime utility directly. Conditional on (n_1, n_2, e_1), expected utility is given by

$$\mathbb{E}[U \mid n_1, n_2, e_1, s_0] = \log\big(w_1(s_0)(1-e_1-\phi n_1)\big) + (1-\pi n_2)\log(c_{2,0}) + \pi n_2 \log(c_{2,1}) + v n_1 + \pi v n_2, \tag{G.13}$$

Career Investment

Given n_1 and s_0, the woman chooses career investment e_1 in period 1 to maximize (G.13). The first-order condition for an interior solution is

$$-\frac{1}{1-e_1-\phi n_1} + \frac{\gamma(s_0)}{e_1} = 0, \tag{G.14}$$

which yields the optimal career investment rule

$$e_1^*(s_0, n_1) = \frac{\gamma(s_0)}{1+\gamma(s_0)}(1-\phi n_1). \tag{G.15}$$

Equation (G.15) implies that, conditional on schooling s_0, early fertility scales down career investment proportionally to the reduction in available time:

$$e_1^*(s_0, 1) = (1-\phi)\, e_1^*(s_0, 0). \tag{G.16}$$

Fertility Plans

Women differ in whether fertility timing remains a choice at the start of period 1. Women with $n_0 = 1$ enter transition to adulthood having already had a child and therefore face motherhood in period 1 ($n_1 = 1$); for these women, fertility timing is predetermined. In contrast, women with $n_0 = 0$ enter period 1 without children and can choose whether and when to have a child.

For women whose fertility timing remains a choice ($n_0 = 0$), fertility decisions are characterized by comparing expected lifetime utility across alternative fertility paths, taking optimal career investment as given. In particular, women choose among three fertility plans: no children, delayed fertility, and early fertility.

Let U^N, U^L, and U^E denote expected lifetime utility under the no-children, delayed-fertility, and early-fertility paths, respectively. Evaluating expected lifetime utility (G.13) at the optimal career investment choice $e_1 = e_1^*(s_0, n_1)$ yields

$$U^N = \log\left(\frac{w_1(s_0^*(y,0))}{1+\gamma(s_0^*(y,0))}\right) + \log\left(w_1(s_0^*(y,0))\,\kappa(s_0)\left(\frac{\gamma(s_0^*(y,0))}{1+\gamma(s_0^*(y,0))}\right)^{\gamma(s_0^*(y,0))}\right), \tag{G.17}$$

$$U^E = U^N + \big(1+\gamma(s_0^*(y,0))\big)\log(1-\phi) + v, \tag{G.18}$$

$$U^L = U^N + \pi\log(1-\phi) + \pi v. \tag{G.19}$$

The fertility decision can be characterized by two thresholds in the utility value of children v. First, delayed fertility is preferred to childlessness if and only if

$$v \geq \bar{v}_1 \equiv -\log(1-\phi). \tag{G.20}$$

Second, early fertility is preferred to delayed fertility if and only if

$$v \geq \bar{v}_2\big(s_0^*(y,0)\big) \equiv \frac{1+\gamma(s_0^*(y,0))-\pi}{1-\pi}\big[-\log(1-\phi)\big]. \tag{G.21}$$

Combining these conditions, the optimal fertility path for women with $n_0 = 0$ is

$$(n_1, n_2) = \begin{cases} (0,0) & \text{if } v < \bar{v}_1 \quad \text{(no children)}, \\ (0,1) & \text{if } \bar{v}_1 \leq v < \bar{v}_2\big(s_0^*(y,0)\big) \quad \text{(delayed fertility)}, \\ (1,0) & \text{if } v \geq \bar{v}_2\big(s_0^*(y,0)\big) \quad \text{(early fertility)}. \end{cases} \tag{G.22}$$

Figure G.1 provides a graphical representation of the resulting choice regions along the support of v. Relative to the original model, the structure of the thresholds is unchanged, but their prevalence now depends on childhood conditions through realized schooling for women whose fertility timing remains a choice, $s_0^*(y,0)$.

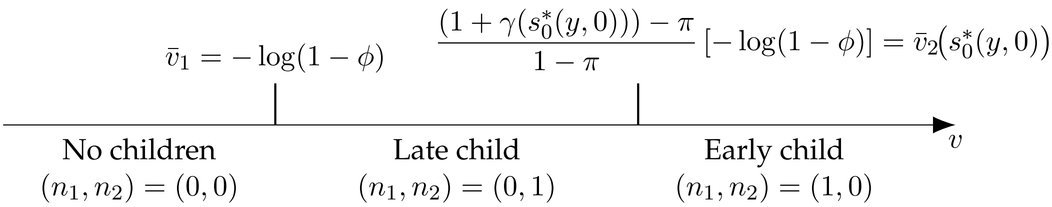

Figure G.1: Type thresholds in v and fertility timing.

Let $F(v)$ denote the cumulative distribution function of v. Given childhood income y, the population shares of women following each fertility path are

$$\omega^E(y) = r(y) + \big(1 - r(y)\big)\Big[1 - F\big(\bar{v}_2(s_0^*(y,0))\big)\Big], \tag{G.23}$$

$$\omega^L(y) = \big(1 - r(y)\big)\Big[F\big(\bar{v}_2(s_0^*(y,0))\big) - F(\bar{v}_1)\Big], \tag{G.24}$$

$$\omega^N(y) = \big(1 - r(y)\big)F(\bar{v}_1). \tag{G.25}$$

In these expressions, the term $r(y)$ captures women who experience a teenage birth and therefore do not make a fertility timing choice, while the factor $1 - r(y)$ reflects women whose fertility timing remains a choice and whose behavior is governed by the thresholds $\bar{v}_1$ and $\bar{v}_2(s_0^*(y,0))$. Importantly, the shares $\omega^E(y)$, $\omega^L(y)$, and $\omega^N(y)$ describe the distribution of optimal fertility plans rather than realized births. In the delayed-fertility group, only a fraction π will realize a birth in period 2. Because π is exogenous and does not vary with childhood income, the distinction between planned and realized fertility does not affect the comparative statics with respect to y discussed below.

Comparative statics with respect to childhood income

Schooling and initial conditions

Consider a positive income shock during childhood. By assumption, the probability of an teenage birth satisfies $r'(y) < 0$. Through this channel, additional childhood income directly reduces the likelihood of entering transition to adulthood having already experienced a birth.

Regarding schooling investment, recall that the optimal period-0 schooling choice is $s_0^*(y, n_0)$ (Equation (G.6)). Holding the early-birth realization n_0 fixed, the partial derivative with respect to childhood income is

$$\frac{\partial s_0^*(y, n_0)}{\partial y} = \frac{\delta}{1+\delta} \cdot \frac{1}{p(1+\varphi n_0)} > 0. \tag{G.26}$$

Thus, conditional on n_0, higher childhood income mechanically increases optimal schooling investment, reflecting greater available resources under the assumption that schooling is a normal good.

While schooling is realized at the individual level as $s_0 = s_0^*(y, n_0)$, it is useful for comparative statics to summarize schooling outcomes at the population level. Define expected schooling at the start of adulthood as

$$\mathbb{E}[s_0 \mid y] \equiv (1 - r(y))\, s_0^*(y, 0) + r(y)\, s_0^*(y, 1),$$

which averages realized schooling across early-birth realizations.

Using this definition, the total derivative of expected schooling with respect to childhood income is

$$\begin{aligned} \frac{d}{dy}\, \mathbb{E}[s_0 \mid y] &= \frac{d}{dy}\Big((1 - r(y))\, s_0^*(y, 0) + r(y)\, s_0^*(y, 1)\Big) \\ &= (1 - r(y))\, \frac{\partial s_0^*(y, 0)}{\partial y} + r(y)\, \frac{\partial s_0^*(y, 1)}{\partial y} + r'(y)\Big[s_0^*(y, 1) - s_0^*(y, 0)\Big] > 0 \qquad \text{(G.27)} \end{aligned}$$

The first two terms capture within-state responses of schooling to higher income. Both are strictly positive by (G.26). The last term captures a composition effect operating through changes in the probability of an early birth before adulthood. Since an early birth raises the effective cost of schooling, $s_0^*(y, 1) < s_0^*(y, 0)$ (Equation (G.7)), and since $r'(y) < 0$, the bracketed term is negative while $r'(y)$ is negative, implying that the product

$$r'(y)\Big[s_0^*(y, 1) - s_0^*(y, 0)\Big] > 0.$$

Therefore, higher childhood income increases expected schooling both by raising schooling investment within each early-birth state and by shifting individuals away from the high-cost state $n_0 = 1$. This increase in expected schooling shapes entry wages and the returns to career investment in period 1.

Fertility timing and realized fertility

Given childhood income y, the population shares of women following each fertility path are given by (G.23)–(G.25). Differentiating the shares with respect to childhood income yields

$$\begin{aligned} \frac{d\,\omega^E(y)}{dy} &= r'(y)\, F(\bar{v}_2) \\ &\quad - (1 - r(y))\, f(\bar{v}_2)\, \bar{v}_2'\, \frac{\partial s_0^*(y, 0)}{\partial y} < 0, \qquad \text{(G.28)} \\ \frac{d\,\omega^L(y)}{dy} &= -r'(y)\Big[F(\bar{v}_2) - F(\bar{v}_1)\Big] \\ &\quad + (1 - r(y))\, f(\bar{v}_2)\, \bar{v}_2'\, \frac{\partial s_0^*(y, 0)}{\partial y} > 0, \qquad \text{(G.29)} \\ \frac{d\,\omega^N(y)}{dy} &= -r'(y)\, F(\bar{v}_1) > 0. \qquad \text{(G.30)} \end{aligned}$$

These expressions highlight two distinct channels through which childhood income affects fertility timing.

First, there is a *mechanical* selection channel operating through teenage birth risk. Since $r'(y) < 0$, higher childhood income reduces the probability of entering transition to adulthood having already had a birth ($n_0 = 1$). This mechanically removes women from the "forced" teenage-fertility group and reallocates them across the three fertility paths according to their preferences. Among these newly unconstrained women, only a fraction $1 - F(\bar{v}_2)$ would choose to stay in the early fertility (period 1) path, while the remaining fraction $F(\bar{v}_2)$ would choose delayed fertility (period 2) or childlessness, as reflected by the first terms in Equations (G.29) and (G.30). Therefore, even if $\bar{v}_2$ remained unchanged, the early-fertility share $\omega^E(y)$ would mechanically decline.

Second, there is a *behavioral* channel operating through schooling among women whose fertility is a choice ($n_0 = 0$). Higher childhood income raises $s_0^*(y, 0)$, which increases the opportunity cost of early fertility by strengthening the incentive to allocate time toward career investment. Formally, because $\gamma'(s_0) > 0$ implies $\bar{v}_2'(s_0) > 0$, the threshold separating early from delayed fertility, $\bar{v}_2(s_0^*(y, 0))$, increases with childhood income, shifting women from early to delayed fertility. In contrast, the threshold separating delayed fertility from childlessness, $\bar{v}_1$, is independent of childhood income, so behavioral responses operate only on the timing margin (E vs L).

Taken together, both channels operate in the same direction for early fertility. The mechanical reduction in early births before adulthood and the behavioral increase in the threshold $\bar{v}_2$ both lower the share of women on the early-fertility path, implying $\frac{d\omega^E(y)}{dy} < 0$. The delayed-fertility share increases through both channels: fewer women are constrained by period-0 fertility, and among women whose fertility timing remains a choice, higher schooling shifts the early-versus-delayed threshold upward. By contrast, the behavioral channel does not directly affect the planned childlessness margin, since $\bar{v}_1$ is independent of childhood income. Changes in planned childlessness therefore arise only through the mechanical reduction in period-0 fertility risk: some women who choose the delayed-fertility path do not realize a birth due to infertility risk. Thus, higher childhood income unambiguously shifts fertility away from early motherhood and toward delayed fertility, while effects on childlessness would arise only through the mechanical reduction in period-0 fertility risk and through unrealized delayed births.

Career investment

To understand how career investment responds to childhood income, it is useful to first consider how it varies with schooling. The partial derivative of optimal career investment with respect to schooling is

$$\frac{\partial e_1^*(s_0, n_1)}{\partial s_0} = \frac{\gamma'(s_0)}{(1 + \gamma(s_0))^2}(1 - \phi n_1) > 0,$$

so higher schooling leads to higher career investment for any fertility realization.

Because schooling increases with childhood income (Equation (G.26)), career investment also increases with childhood income through this channel:

$$\frac{de_1^*}{dy} = \frac{\partial e_1^*(s_0, n_1)}{\partial s_0} \frac{\partial s_0^*(y, n_0)}{\partial y} > 0. \tag{G.31}$$

As for schooling choices, it is useful for comparative statics to summarize career investment at the population level. Given childhood income y, there are three relevant optimal investment levels associated with different fertility realizations:

$$e_1^*\big(s_0^*(y,0),0\big), \qquad e_1^*\big(s_0^*(y,0),1\big), \qquad e_1^*\big(s_0^*(y,1),1\big).$$

The first two correspond to women who enter period 1 without children ($n_0 = 0$), for whom fertility timing remains a choice, while the third corresponds to women who already had a child ($n_0 = 1$). For convenience, define

$$\mu_1(y) \equiv \mathbb{E}[e_1^* \mid y, n_0 = 1] = e_1^*\big(s_0^*(y,1),1\big),$$

$$\mu_0(y) \equiv \mathbb{E}[e_1^* \mid y, n_0 = 0] = F(\bar{v}_2)\, e_1^*\big(s_0^*(y,0),0\big) + \big[1 - F(\bar{v}_2)\big]\, e_1^*\big(s_0^*(y,0),1\big).$$

Expected career investment is therefore

$$\mathbb{E}[e_1^* \mid y] = (1 - r(y))\, \mu_0(y) + r(y)\, \mu_1(y). \tag{G.32}$$

Differentiating (G.32) yields

$$\frac{d}{dy}\mathbb{E}[e_1^* \mid y] = (1 - r(y))\, \mu_0'(y) + r(y)\, \mu_1'(y) + r'(y)\big(\mu_1(y) - \mu_0(y)\big). \tag{G.33}$$

where

$$\begin{aligned} \mu_0'(y) = {} & f(\bar{v}_2)\, \bar{v}_2' \, \frac{\partial s_0^*(y,0)}{\partial y} \left[e_1^*\big(s_0^*(y,0),0\big) - e_1^*\big(s_0^*(y,0),1\big) \right] \\ & + F(\bar{v}_2)\, \frac{d}{dy} e_1^*\big(s_0^*(y,0),0\big) + \big[1 - F(\bar{v}_2)\big]\, \frac{d}{dy} e_1^*\big(s_0^*(y,0),1\big) > 0, \end{aligned} \tag{G.34}$$

and

$$\mu_1'(y) = \frac{d}{dy}\, e_1^*\big(s_0^*(y,1),1\big) > 0. \tag{G.35}$$

The first two terms in Equation (G.33) capture changes in career investment within each period-0 fertility state. In particular, $\mu_0'(y)$ reflects changes among women with $n_0 = 0$, for whom fertility timing remains a choice in period 1; therefore, it includes both the direct effect of higher childhood income on career investment through schooling and a composition effect within this group arising from shifts between early and delayed fertility as $\bar{v}_2$ changes. By contrast, $\mu_1'(y)$ captures changes among women with $n_0 = 1$, whose fertility timing is predetermined. From previous results, both $\mu_0'(y)$ and $\mu_1'(y)$ are unambiguously positive. Intuitively, for any given fertility realization, higher childhood income increases career investment (Equation G.31). In addition, within the group of women with $n_0 = 0$, higher income shifts the threshold $\bar{v}_2$, reallocating women from early to delayed fertility and further increasing average career investment (Equation G.16).

The last term in Equation (G.33) can instead be interpreted as changes in career investment associated with shifts in period-0 fertility states. From Equations (G.16) and (G.7), career investment is lower under early fertility than under delayed or no fertility due to the time cost ϕ, and optimal schooling is lower when period-0 fertility occurs. Because career investment is increasing in schooling, we can establish the following ranking:

$$e_1^*\big(s_0^*(y,1),1\big) < e_1^*\big(s_0^*(y,0),1\big) < e_1^*\big(s_0^*(y,0),0\big).$$

This implies that

$$\mu_1(y) - \mu_0(y) < 0.$$

Furthermore, since $r'(y) < 0$ by assumption, reflecting the negative relation between childhood income and period-0 fertility, the term $r'(y)\big(\mu_1(y) - \mu_0(y)\big)$ is positive. Taken together, all three terms in Equation (G.33) are positive. Hence, higher childhood income unambiguously increases career investment in period 1.

Period 1 earnings

Multiplying wages by the time allocated to work, period 1 earnings can be written as

$$Y_1(s_0, n_1) = w_1(s_0)\,(1 - e_1 - \phi n_1). \tag{G.36}$$

Substituting the optimal career investment rule from Equation (G.15) yields

$$Y_1(s_0, n_1) = w_1(s_0)\,\frac{1-\phi n_1}{1+\gamma(s_0)}. \tag{G.37}$$

This expression highlights that period 1 earnings reflect both productivity and time allocation decisions. In particular, $w_1(s_0)$ captures productivity through entry wages, while the term $\frac{1}{1+\gamma(s_0)}$ reflects the endogenous allocation of time between work and career investment, and $(1-\phi n_1)$ captures the time cost of fertility.

Differentiating with respect to s_0,

$$\frac{\partial Y_1(s_0,n_1)}{\partial s_0} = \frac{1-\phi n_1}{1+\gamma(s_0)}\left[w_1'(s_0) - w_1(s_0)\frac{\gamma'(s_0)}{1+\gamma(s_0)}\right].$$

The sign of this derivative is ambiguous. Higher schooling raises earnings through higher entry wages, but it also increases the returns to career investment, which induces higher investment and reduces time allocated to work. As a result, the effect of schooling on period 1 earnings reflects a tradeoff between increased entry wages and reduced labor supply due to endogenous career investment. Hence,

$$\frac{dY_1(s_0,n_1)}{dy} = \frac{\partial Y_1(s_0,n_1)}{\partial s_0}\frac{\partial s_0^*(y,n_0)}{\partial y}$$

is also ambiguous.

As in the case of career investment, it is useful to decompose expected period 1 earnings by period-0 fertility state. Let

$$\nu_1(y) \equiv \mathbb{E}[Y_1 \mid y, n_0 = 1] = Y_1\big(s_0^*(y,1),1\big),$$

$$\nu_0(y) \equiv \mathbb{E}[Y_1 \mid y, n_0 = 0] = F(\bar{v}_2)\,Y_1\big(s_0^*(y,0),0\big) + \big[1 - F(\bar{v}_2)\big]\,Y_1\big(s_0^*(y,0),1\big).$$

Expected period 1 earnings are therefore

$$\mathbb{E}[Y_1 \mid y] = (1 - r(y))\, \nu_0(y) + r(y)\, \nu_1(y). \tag{G.38}$$

Differentiating (G.38) yields

$$\frac{d}{dy}\mathbb{E}[Y_1 \mid y] = (1 - r(y))\, \nu_0'(y) + r(y)\, \nu_1'(y) + r'(y)\big(\nu_1(y) - \nu_0(y)\big). \tag{G.39}$$

where

$$\begin{aligned} \nu_0'(y) = f(\bar{v}_2)\, \bar{v}_2' \,\frac{\partial s_0^*(y,0)}{\partial y} \left[Y_1\big(s_0^*(y,0),0\big) - Y_1\big(s_0^*(y,0),1\big)\right] \\ + F(\bar{v}_2)\, \frac{d}{dy} Y_1\big(s_0^*(y,0),0\big) + \left[1 - F(\bar{v}_2)\right] \frac{d}{dy} Y_1\big(s_0^*(y,0),1\big), \end{aligned} \tag{G.40}$$

and

$$\nu_1'(y) = \frac{d}{dy} Y_1\big(s_0^*(y,1),1\big). \tag{G.41}$$

The first two terms in Equation (G.39) capture changes in period 1 earnings within each period-0 fertility state. In particular, these terms combine changes in productivity through $w_1(s_0)$ with changes in time allocation driven by both career investment and fertility. Within the group of women with $n_0 = 0$, the composition effect associated with the first line in Equation (G.40) is unambiguously positive, as shifting women from early to delayed fertility relaxes the period 1 time constraint. The additional ϕ units of time are then optimally allocated between work and career investment, so that a fraction is devoted to work, increasing contemporaneous earnings.

However, the remaining components of $\nu_0'(y)$ and $\nu_1'(y)$ are generally ambiguous, reflecting the endogenous adjustment of career investment. Higher schooling increases the returns to investment, leading individuals to reallocate time away from work, so that gains in productivity need not translate into higher period 1 earnings.

The last term in Equation (G.39) captures changes in period 1 earnings associated with shifts in period-0 fertility states. In particular, it reflects the difference between outcomes for women with $n_0 = 1$ and those with $n_0 = 0$, who differ in both schooling and fertility timing. As a result, this term compares earnings for women with lower schooling and early fertility, $Y_1(s_0^*(y,1),1)$, to those for women with higher schooling and more flexible fertility choices, $Y_1(s_0^*(y,0),0)$ and $Y_1(s_0^*(y,0),1)$. Its sign is therefore ambiguous and depends on the same tradeoff between higher wages and increased career investment discussed above.

Taken together, the effect of childhood income on period 1 earnings reflects a tradeoff between productivity gains and reduced labor supply due to career investment. A sufficient condition for the effect to be positive is

$$w_1'(s_0) > w_1(s_0)\frac{\gamma'(s_0)}{1 + \gamma(s_0)}, \tag{G.42}$$

under which higher schooling translates into higher Y_1. When dynamic complementarities are strong, however, individuals devote a larger share of time to career investment, so that period 1 earnings may respond weakly or even decline despite higher schooling.

Period 2 earnings

Using the optimal career investment rule from Equation (G.15), period 2 earnings can be written as

$$Y_2(s_0, n_1, n_2) = w_1(s_0)\,\kappa(s_0)\,\left[\frac{\gamma(s_0)}{1+\gamma(s_0)}(1-\phi n_1)\right]^{\gamma(s_0)}(1-\phi n_2). \tag{G.43}$$

It is useful to separate period 2 productivity/wages from realized earnings. Define

$$w_2(s_0, n_1) = w_1(s_0)\,\kappa(s_0)\,\left[\frac{\gamma(s_0)}{1+\gamma(s_0)}(1-\phi n_1)\right]^{\gamma(s_0)},$$

so that $Y_2(s_0, n_1, n_2) = w_2(s_0, n_1)(1-\phi n_2)$.

Under the assumption introduced above, which ensures that the increase in the productivity scale $\kappa(s_0)$ dominates the negative mechanical effect of $\gamma(s_0)$ on the wage multiplier due to $e \in (0,1)$, higher schooling raises period 2 wages through higher entry wages, a higher productivity scale of career investment, and stronger incentives for career investment. This implies

$$\frac{\partial w_2}{\partial s_0} > 0,$$

and therefore

$$\frac{dw_2}{dy} = \frac{\partial w_2}{\partial s_0}\frac{\partial s_0^*(y, n_0)}{\partial y} > 0.$$

Importantly, this wage channel does not depend on n_2. The role of ϕ in period 2 earnings arises instead through the multiplicative term $(1-\phi n_2)$, which scales realized earnings.

As in the previous cases, it is useful to summarize period 2 earnings at the population level. Given childhood income y, define

$$\eta_1(y) \equiv \mathbb{E}[Y_2 \mid y, n_0 = 1] = Y_2\big(s_0^*(y,1), 1, 0\big),$$

$$\begin{aligned}\eta_0(y) \equiv \mathbb{E}[Y_2 \mid y, n_0 = 0] = {} & F(\bar{v}_1)\,Y_2\big(s_0^*(y,0), 0, 0\big)\\ & + \big[F(\bar{v}_2) - F(\bar{v}_1)\big]\,Y_2\big(s_0^*(y,0), 0, 1\big)\\ & + \big[1 - F(\bar{v}_2)\big]\,Y_2\big(s_0^*(y,0), 1, 0\big).\end{aligned}$$

Expected period 2 earnings are therefore

$$\mathbb{E}[Y_2 \mid y] = (1 - r(y))\,\eta_0(y) + r(y)\,\eta_1(y). \tag{G.44}$$

Differentiating (G.44) yields

$$\frac{d}{dy}\mathbb{E}[Y_2 \mid y] = (1 - r(y))\,\eta_0'(y) + r(y)\,\eta_1'(y) + r'(y)\big(\eta_1(y) - \eta_0(y)\big). \tag{G.45}$$

where

$$\eta_0'(y) = f(\bar{v}_2)\,\bar{v}_2'\,\frac{\partial s_0^*(y,0)}{\partial y}\,\big[Y_2\big(s_0^*(y,0),0,1\big) - Y_2\big(s_0^*(y,0),1,0\big)\big]$$
$$+ F(\bar{v}_1)\,\frac{d}{dy}Y_2\big(s_0^*(y,0),0,0\big) + \big[F(\bar{v}_2) - F(\bar{v}_1)\big]\,\frac{d}{dy}Y_2\big(s_0^*(y,0),0,1\big)$$
$$+ \big[1 - F(\bar{v}_2)\big]\,\frac{d}{dy}Y_2\big(s_0^*(y,0),1,0\big), \tag{G.46}$$

and

$$\eta_1'(y) = \frac{d}{dy}\,Y_2\big(s_0^*(y,1),1,0\big). \tag{G.47}$$

The first two terms in Equation (G.45) capture changes in period 2 earnings within each period-0 fertility state. From our previous assumptions, the terms associated with $\frac{d}{dy}Y_2(\cdot)$ are unambiguously positive, as $Y_2 = w_2(1-\phi n_2)$. However, it is important to note that changes in earnings associated with wages/productivity gains are always scaled down by the factor $(1-\phi)$ whenever $n_2 = 1$.

In any case, the sign of the composition effect within $\eta_0'(y)$ still depends on the comparison between delayed and early fertility for a given level of schooling,

$$Y_2\big(s_0^*(y,0),0,1\big) - Y_2\big(s_0^*(y,0),1,0\big),$$

which can be expressed as a positive factor multiplied by

$$(1-\phi) - (1-\phi)^{\gamma\left(s_0^*(y,0)\right)}.$$

Since $1-\phi \in (0,1)$, delayed fertility yields higher period 2 earnings than early fertility if and only if $\gamma(s_0^*(y,0)) > 1$. Intuitively, early fertility reduces career investment in period 1, while delayed fertility instead reduces contemporaneous earnings in period 2 through the time cost ϕ. When the elasticity of period-2 wages with respect to career investment is sufficiently strong, the loss from reduced investment dominates the contemporaneous time cost, so delayed fertility yields higher period 2 earnings.

The last term in Equation (G.45) captures changes in period 2 earnings associated with shifts in period-0 fertility states. Using the fact that period 2 wages are increasing in schooling and career investment, while realized earnings are decreasing in n_2, we can rank outcomes as

$$Y_2\big(s_0^*(y,0),0,0\big) > Y_2\big(s_0^*(y,0),0,1\big) > Y_2\big(s_0^*(y,0),1,0\big) > Y_2\big(s_0^*(y,1),1,0\big),$$

under the condition $\gamma(s_0^*(y,0)) > 1$, which implies $\eta_1(y) < \eta_0(y)$ and therefore a positive contribution from this selection channel.

Taken together, these results show that higher childhood income increases period 2 wages through higher schooling, a higher productivity scale of career investment, and greater career investment. However, because realized earnings are scaled by $(1-\phi n_2)$, part of these gains may not be reflected in observed earnings when fertility is delayed. As childhood income increases and more women shift from early to delayed fertility, the incidence of the child penalty moves from period 1 to period 2. Consequently, the effect of childhood income on period 2 earnings may be attenuated. The extent of attenuation increases with ϕ, as higher time costs reduce the fraction of productivity that is translated into realized earnings. For instance, in the extreme case where motherhood and employment are mutually exclusive, productivity gains are not translated into earnings at all because there is no labor market participation, at least temporarily.

Discussion: Earnings over the life cycle

Period 1 earnings, period 2 earnings, and cumulative earnings are shaped by the same underlying channels: higher childhood income increases schooling, reduces exposure to early fertility, and raises career investment. At the same time, earnings at a given age reflect how individuals allocate time between work, career investment, and childbearing, so that multiple forces, some increasing and some reducing labor supply, operate simultaneously, and realized earnings need not move one-for-one with underlying productivity.

In period 1, changes in earnings reflect both changes in productivity and the allocation of time between work, career investment, and childbearing. Three forces operate simultaneously: higher entry wages through higher s_0, a relaxation of time constraints as early fertility declines, and a reallocation of time toward career investment as its returns increase. From the previous section, a sufficient condition for period 1 earnings to increase with childhood income is

$$w_1'(s_0) > w_1(s_0)\frac{\gamma'(s_0)}{1+\gamma(s_0)},$$

which requires the direct effect of schooling on wages to dominate the induced increase in career investment. In addition, shifts away from early fertility relax time constraints and increase the time available for work, contributing positively to earnings. Taken together, these forces suggest that period 1 earnings are most likely to increase with childhood income. However, when dynamic complementarities are strong and the effects of y on s_0 are large, higher schooling leads to substantial reallocation of time toward career investment, so that gains in productivity need not translate into higher contemporaneous earnings. In extreme cases, this may result in weak or even negative earnings responses if individuals temporarily withdraw from the labor market to make strong career investments.

In period 2, higher childhood income raises wages/productivity through higher entry wages, a higher productivity scale of career investment, greater career investment, and stronger returns to that investment, while realized earnings depend on the timing of fertility through its effect on the time available for work. In particular, women who delay fertility realize the time cost of childbearing in period 2, which attenuates the translation of productivity gains into observed earnings. As a result, the wage/productivity component of period 2 earnings is expected to increase with childhood income, but the magnitude of these effects in terms of actual earnings may be limited in settings where employment and motherhood are strong substitutes.

To summarize the overall impact on earnings over the life cycle, define cumulative earnings as

$$Y^T = Y_1 + Y_2.$$

Childhood income affects cumulative earnings through persistent improvements in schooling and career investment, as well as through a reallocation of fertility over the life cycle. While earnings at a given age may be reduced by increased career investment in period 1 or by the realization of fertility costs in period 2, these forces reflect the timing of decisions rather than permanent reductions in productivity. By contrast, the wage/productivity gains generated by higher schooling and career investment are structural and persist over time. As a result, the model predicts more clearly positive effects on cumulative earnings, even when earnings at specific ages are attenuated. Over longer horizons, as career investments pay off and fertility timing adjusts, these temporary effects become less important, leading to positive impacts on overall earnings.

H Marginal Value of Public Funds Calculations

This appendix describes in detail the construction of the Marginal Value of Public Funds (MVPF) estimates reported in Section 5.4. Although the empirical analysis follows individuals through 2023, the effects of childhood income support may continue to accumulate beyond the observed period. I therefore proceed in two steps. First, I compute the short-run MVPF using outcomes observed up to 2023, relying exclusively on estimated causal effects and thus requiring minimal additional assumptions. These results are reported in Table H.1. Second, I project the MVPF through retirement age under alternative scenarios regarding the life-cycle evolution of formal labor earnings.

H.1 Short-Run MVPF

Household Characteristics To remain consistent with the empirical design, I report the main MVPF estimates for the average household in the *main sample*, defined as households located within 2 percentage points of the eligibility threshold. All MVPF calculations are expressed at the household level, as this is the level at which the transfer is made. In terms of household characteristics, this average household comprises 4.3 members, of whom 2.4 are children and 1.9 are adults. Among children, 49% are girls. Hence, the expected number of girls per household is $2.4 \times 0.49 = 1.176$. Appendix C provides more detailed descriptive statistics.

Time Discounting Factor. I start by computing the MVPF of an additional USD 1 in *PANES/AFAM-PE* using cumulative behavioral responses observed up to 2023. For simplicity, I do not compute year-by-year effects. Instead, I take the cumulative effects and apply midpoint discounting, i.e., 9 years, corresponding to the midpoint of the 18-year period between 2005 and 2023. Under the standard assumption of a 3% annual discount rate, the discount factor applied to cumulative effects is $(1/1.03)^9$.

Household Willingness to Pay

Willingness to Pay – Adults Since *PANES/AFAM-PE* consists of a cash transfer with almost no restrictions on how it can be spent, adults are assumed to value the transfer at face value. Hence, adults' willingness to pay for an additional USD 1 transferred to the household equals 1.

Willingness to Pay – Children Children's willingness to pay (WTP) is defined as the discounted after-tax increase in earnings induced by the transfer. As in Bergolo and Cruces (2021), let $\tau = 0.32$ denote the average tax rate in Uruguay. This rate includes employee and employer payroll taxes as well as mandatory health insurance contributions, and assumes no personal income taxation. This assumption is reasonable given that in Uruguay the bottom 70% of the labor income distribution is exempt from personal income tax. From Table 2, column (4), the effect of an additional USD 1,000 in childhood income support on girls' cumulative adult labor earnings is 1.740 (in thousands of USD). Hence, the individual-level willingness to pay of a girl is computed as:

$$1.740 \times (1 - 0.32) \times (1/1.03)^9.$$

Because Table 2, column (8), shows that the effect on boys' adult earnings is statistically indistinguishable from zero, household-level children's WTP is obtained by scaling girls' WTP by the expected number of girls per household:

$$\left[1.740 \times 0.68 \times (1/1.03)^9\right] \times 1.176 = 1.066.$$

Willingness to Pay - Household Total household willingness to pay is therefore:

$$1 \; + \; 1.066 \; = \; 2.066.$$

Fiscal Externalities – Adults In the baseline specification, I focus on first-order fiscal effects operating through changes in payroll tax revenues driven by formal labor earnings. On the adult side, the income eligibility thresholds used in *PANES/AFAM-PE* may themselves have generated behavioral responses in parental labor supply, as illustrated in Figure 1 and documented in Bergolo and Cruces (2021). To account for these potential responses, I replicate the main estimation strategy described in Equation 1, using parental labor earnings as the outcome variable. For consistency with the child-level analysis, I use the same endogenous variable in the parental regressions.

In practice, I measure cumulative earnings for the household head and partner in the main sample, and estimate child–parent level regressions weighted by the inverse of the number of children in the household. This weighting prevents households with more children from receiving disproportionate weight in the estimation. The weighted fuzzy RDD estimate indicates an effect of USD -383 on cumulative parental earnings. Applying the payroll tax rate, multiplying by the number of adults in the household, and discounting yields:

$$0.383 \times 0.32 \times 1.9 \times (1/1.03)^9 = 0.178.$$

This represents a reduction in payroll tax revenues per USD 1 transferred.

Fiscal Externalities – Children On the children's side, the increase in formal earnings generates additional payroll tax revenues equal to

$$1{,}740 \times 0.32 \times 1.176 \times (1/1.03)^9 = 0.502.$$

Since this represents an increase in government revenue, it enters the fiscal externality component with a negative sign.

Total fiscal externalities Total fiscal externalities therefore equal:

$$0.196 \; - \; 0.530 \; = \; -0.323.$$

Static MVPF The static MVPF, reported in Table H.1, column (1), is defined as

$$MVPF = \frac{WTP^{HH}}{1 + FE^{HH}},$$

where 1 denotes the mechanical cost of transferring one dollar and FE^{HH} denotes total fiscal externalities per dollar transferred. Substituting the values above yields

$$\frac{2.066}{1 - 0.334} = 3.1.$$

Thus, each additional dollar of net government spending generates approximately 3.1 dollars of benefits for beneficiary households.

Sensitivity: Real vs. Shifting Responses and Additional Fiscal Externalities. So far, I have assumed that effects on formal earnings are entirely driven by real labor supply responses. However, part of the observed response may reflect shifting between formal and informal earnings. This distinction is relevant for MVPF calculations, as not every additional dollar earned in the formal sector necessarily implies an equivalent increase in total income. To assess sensitivity to the assumption that all observed earnings responses correspond to real increases in total income, I allow only a fraction of the estimated earnings effects to reflect real income changes. I also allow for additional fiscal externalities operating through other tax bases, such as VAT revenues, via changes in consumption. I assume that individuals value labor market earnings only through consumption; that is, there is no differential value attached to formal versus informal earnings, and shifts between them do not generate changes in willingness to pay. Let $c = 0.8$ denote the marginal propensity to consume out of after-tax income, and let the VAT rate be 0.22.

In the most conservative scenario, reported in Table H.1, column (2), I assume that all adult earnings responses are real, while all children's earnings responses reflect shifting between formal and informal income.

In this case, adults' willingness to pay remains equal to the face value of the transfer, while children's willingness to pay is set to zero, as earnings are assumed to be valued only through consumption (i.e., there is no intrinsic value attached to formal earnings relative to informal earnings).

Adults' fiscal externalities correspond to the reduction in payroll tax revenues discussed above, plus an additional term associated with changes in VAT revenues due to reduced consumption. Since all adult earnings responses are assumed to be real, the reduction in after-tax income equals

$$383 \times 0.68 \times 1.9$$

and the associated discounted VAT revenue loss equals

$$383 \times 0.68 \times 1.9 \times 0.8 \times 0.22 \times (1/1.03)^9.$$

On the children's side, since all earnings responses are assumed to reflect shifting, VAT externalities are zero because total income and consumption do not change. Children's fiscal externalities therefore correspond only to changes in payroll tax revenues associated with the increase in formal income, as in the baseline scenario. Under these assumptions, the MVPF equals 1.3.

In the intermediate scenario, reported in Table H.1, column (3), I assume that 50% of both adult and child earnings responses are real. In this case, adults' willingness to pay remains unchanged, while children's willingness to pay is reduced to 50% of the after-tax cumulative earnings effects. Fiscal externalities in payroll taxes remain as in the baseline specification. However, both adults and children now generate fiscal externalities through the VAT base, which are computed analogously to the conservative scenario, applying the marginal propensity to consume and the VAT rate to the real component of after-tax income changes. Under this parameterization, the MVPF equals 2.6.

Finally, in the most optimistic scenario, reported in Table H.1, column (4), I assume that all child earnings responses are real and all adult responses reflect shifting between formal and informal income. In this case, VAT revenues increase due to higher consumption by children, while no VAT losses arise from adults. The resulting MVPF equals 4.2.

Across all specifications, household willingness to pay exceeds the fiscal cost of the program.

H.2 Long-Run MVPF

Next, I describe the procedure used to construct the longer-run MVPF estimates. The approach follows the methodology in Hendren and Sprung-Keyser (2020) and Bailey et al. (2024), in which observed earnings effects are extrapolated over the life cycle using age-income profiles from external data sources, such as household surveys. For simplicity, in this long-term analysis I focus on projecting the future stream of fiscal externalities associated with behavioral responses in labor market outcomes beyond age 28, the average last age at which outcomes are directly observed in the administrative data. This allows me to examine when projected fiscal externalities are expected to offset the direct cost of an additional USD 1 transferred through *PANES/AFAM-PE* and to estimate the corresponding fiscal gains at retirement age.

The projection proceeds in three steps. First, I construct a cross-sectional age-earnings profile using nationally representative survey data. Second, I scale this profile to match the earnings level of *PANES/AFAM-PE* ineligible women in my estimation sample. Third, I impute a permanent proportional earnings gap equal to the estimated 2SLS last-year treatment effect and project earnings forward under an assumed wage growth rate. The implied stream of payroll tax revenues is then discounted and accumulated over time.

Step 1: Cross-sectional age-earnings profile from the ECH. I use the 2023 National Household Survey (Encuesta Continua de Hogares, ECH) to construct a cross-sectional earnings profile for women aged 15 to 68. The sample is restricted to women and to working-age individuals. Monthly labor income is annualized by multiplying by 13 to account for the mandatory "aguinaldo" (or yearly bonus) payment in Uruguay. To maintain consistency with the administrative data, I set labor income equal to zero for informal workers (those not contributing to the pension system). I then compute the average annual labor income at each age $a \in \{15, \dots, 68\}$, denoted $\overline{Y}_a^{ECH}$. This object captures the cross-sectional age profile of formal labor earnings in the Uruguayan economy in 2023.

Step 2: Aligning the ECH profile to the AFAM sample. The ECH profile reflects the full national population, while the relevant counterfactual in the MVPF calculation is the earnings path of *PANES/AFAM-PE* ineligible women in my estimation sample. To align levels, I compute the average last-year formal earnings of ineligible women in my sample for ages 21 to 30 and compare them to the corresponding ECH averages.

Let $\overline{Y}_a^{inelig}$ denote average earnings of ineligible women in the administrative data at age a, and let $\overline{Y}_a^{ECH}$ denote the corresponding ECH average. I compute the ratio

$$\rho = \text{mean}_a \left(\frac{\overline{Y}_a^{inelig}}{\overline{Y}_a^{ECH}} \right),$$

averaged across observed ages. This ratio captures the relative earnings level of the ineligible AFAM sample compared to the national cross-section.

I then construct a simulated lifecycle earnings profile for ineligible women as

$$Y_a^{inelig} = \rho \times \overline{Y}_a^{ECH} \times (1.005)^{a-28}.$$

The term $(1.005)^{a-15}$ introduces a 0.5% annual real wage growth assumption. Alternative projections with 0.25% and 0.75% growth are also constructed for robustness.

Step 3: Imposing the treatment effect. The dynamic projection assumes that eligible women follow the same lifecycle earnings shape as ineligible women but at a permanently higher level. The proportional gap is pinned down by the 2SLS last-year treatment effect reported in Table 2, column (2), equal to 5.2%. Thus, simulated earnings for eligible women are defined as

$$Y_a^{elig} = Y_a^{inelig} \times (1 + 0.052).$$

Step 4: Projected fiscal externalities. To compute the projected change in fiscal externalities associated with behavioral responses induced by the program, I take the income–age profiles estimated in step 3, compute the difference between the eligible and ineligible groups (i.e., the simulated treatment effect), and then proceed as in the baseline scenario:

$$\left[Y_a^{elig} - Y_a^{inelig}\right] \times 0.32 \times 1.176 \times (1/1.03)^{(a-28)} \times (1/1.03)^9.$$

The final term, $(1/1.03)^9$, aligns projected post-28 revenues with the same present-value reference date used in the short-run calculations, which apply midpoint discounting to cumulative effects observed between 2005 and 2023. The discounted revenues are then summed over ages 29 through retirement age (68) to obtain total additional fiscal externalities beyond age 28. The cumulative projected fiscal externalities are added to the fiscal externalities already observed up to age 28 in the baseline scenario of the short-run MVPF calculations.

Step 5: Simulated decay rates. The specification discussed so far uses a constant proportional earnings gap over the lifecycle, providing a convenient benchmark for extrapolating long-run fiscal effects. However, it may be misleading if effects on last-year earnings decline over time. To assess the sensitivity of the results to alternative persistence assumptions, I also consider scenarios in which the treatment effect decays over time. In these cases, the earnings difference is scaled by a factor $(1-d)^{(a-28)}$, so that the fiscal contribution at age a becomes

$$\left[Y_a^{elig} - Y_a^{inelig}\right] \times (1-d)^{(a-28)} \times 0.32 \times 1.176 \times (1/1.03)^{(a-28)} \times (1/1.03)^9,$$

where $d \in \{0, 0.005, 0.01, 0.025, 0.05, 0.1\}$ corresponds to yearly decay rates ranging from 0 to 10%.

Baseline projected cumulative fiscal externalities by age are reported (in absolute value) in Figure H.1 panel (a), where the horizontal dashed line indicates the point at which fiscal externalities more than compensate for the direct cost of an additional dollar of *PANES/AFAM-PE* transfers, that is, where the MVPF becomes infinite. Because projected payroll tax revenues eventually exceed the mechanical transfer cost, cumulative fiscal externalities become negative (i.e., net revenue gains), implying that the MVPF tends toward infinity. Furthermore, by retirement age, projected fiscal externalities amount to USD 2.4 per dollar transferred (2.3 under a 0.25% wage growth assumption and 2.5 under a 0.75% assumption), implying net revenue gains of USD 1.4 per dollar transferred. Thus, while the static estimates already indicate that beneficiary households' willingness to pay exceeds the program's fiscal cost, the dynamic projections suggest that, under standard assumptions, the program is likely to break even relatively early in the life cycle and to generate net fiscal surpluses thereafter, implying an infinite MVPF. Moreover, panel (b) shows that this conclusion is robust to allowing the earnings gap to decay over time. Even under relatively fast decay rates, cumulative fiscal externalities remain large and continue to exceed program costs over the lifecycle. In particular, this also holds when allowing the earnings gap to decay at rates as high as 10% per year, indicating that the results are not driven by the assumption of a permanent proportional earnings gap.

Table H.1: Marginal Value of Public Funds (MVPF)

		Formal/Informal Earnings Shifting		
	Baseline (1)	Conservative (2)	Moderate (3)	Optimistic (4)
a. Willingness to Pay (per USD 1 transferred)				
Adults	1.000	1.000	1.000	1.000
Children	1.066	0.000	0.533	1.066
Total WTP	2.066	1.000	1.533	2.066
b. Fiscal Externalities (per USD 1 transferred)				
Adults: payroll taxes	0.178	0.178	0.178	0.178
Adults: VAT	.	0.067	0.033	0.000
Children: payroll taxes	-0.502	-0.502	-0.502	-0.502
Children: VAT	.	.	-0.094	-0.188
Total Fiscal Externalities	-0.323	-0.257	-0.384	-0.511
c. Marginal Value of Public Funds				
MVPF	3.054	1.345	2.488	4.226

Notes: This table reports the Marginal Value of Public Funds (MVPF) associated with an additional USD 1 of *PANES/AFAM-PE* transfers, computed for the average household in the main sample in December 2023, when children are on average 28 years old. Panel A reports beneficiaries' willingness to pay (WTP), Panel B reports fiscal externalities, and Panel C combines these components to compute the MVPF as the ratio of total WTP to the net government cost, defined as $\frac{1}{1+FE}$. All calculations are expressed per USD 1 transferred. In column (1), i.e., the baseline scenario, adults are assumed to value transfers at face value. Children's WTP is given by the discounted increase in after-tax formal labor earnings induced by the program. The cumulative earnings effect for children is 1.74, evaluated using an average payroll tax rate of 32% and discounted at an annual rate of 3 percent using a midpoint approximation (nine years). Individual-level estimates are aggregated to the household level assuming 2.4 children per household, of whom 49 percent are girls (1.176 girls per household). Throughout, children's earnings gains are assumed to accrue only to women. Fiscal externalities in the baseline scenario arise exclusively from payroll tax revenues associated with changes in formal labor earnings. For adults, the cumulative earnings effect is -0.383 per adult, applied to 1.9 adults per household and evaluated using the same tax rate and discounting procedure as for children. For children, increased formal earnings generate additional payroll tax revenues, computed analogously. Column (2) reports a conservative scenario in which adults' earnings responses are 100% real while children's earnings responses are assumed to be purely shifting. Furthermore, children's WTP is set to zero so that formal income is valued in exactly the same way as informal income. In addition, value-added tax revenues are included among fiscal externalities, assuming a VAT rate of 22 percent and a marginal propensity to consume of 0.8 out of after-tax earnings. Column (3) reports a moderate scenario in which 50% of both adults' and children's earnings responses are assumed to be real. Column (4) reports an optimistic scenario in which children's earnings responses are fully real and adults' responses are fully shifting. In all columns, adults' WTP remains equal to one, and payroll tax revenues are computed based on formal earnings.

Figure H.1: Dynamics of MVPF

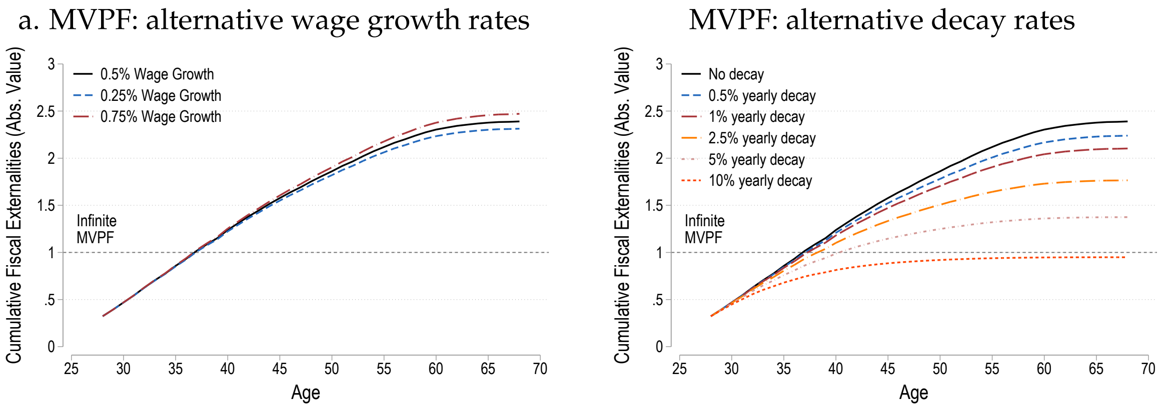

Notes: This figure depicts the evolution of cumulative fiscal externalities over the life cycle based on projected lifetime earnings. The baseline fiscal externality at age 28 corresponds to the static estimate reported in Column (1) of Table H.1. The x-axis reports age, while the y-axis shows the absolute value of cumulative fiscal externalities, discounted at an annual rate of 3 percent. In panel (a), projected earnings paths using the cross-sectional age profile of formal labor earnings observed in the 2023 National Household Survey (ECH), rescaled to match the average income level of ineligible women near the eligibility cutoff, and assuming a 0.5% annual wage growth rate under the baseline scenario. Eligible women are assumed to follow the same lifecycle earnings shape, but at a permanently higher level, with the proportional earnings gap calibrated to the estimated last-year treatment effect reported in Column (2) of Table 2 (i.e., 5.2%). Cumulative fiscal externalities are computed by translating projected earnings gains into additional payroll tax revenues and cumulating these effects over time, including the fiscal effects already observed at age 28. The horizontal reference line at one indicates the level at which cumulative fiscal externalities reach 1, implying that the net fiscal cost of the program is fully offset by future tax revenues and that the Marginal Value of Public Funds becomes infinite. Alternative scenarios using 0.25% and 0.75% wage growth rates are reported by dashed lines. Panel (b) replicates the same exercise allowing the projected earnings gap to decay over time. Specifically, the proportional earnings difference between eligible and ineligible women is assumed to decline at constant annual rates ranging from 0.5% to 10%. All other components of the projection, including the underlying earnings profile and the 0.5% wage growth assumption, remain unchanged.